

# MACTECH

The Journal of Macintosh Technology

# 2006 Sampler

# Welcome to MacTech's 2006 Sampler

Last year's MacTech Sampler was so successful, that we thought we'd do it again! But, this year we've not only given you a wonderful selection of articles, but we're showing you the type of advertisers that we have. This is important since most MacTech readers find the ads useful.

On the following pages, we took a sample of articles we've produced in the last year to give you a good look at what the new MacTech has to offer. There's something for everyone ... network admins to IT Pros, hobbyist to Enterprise, scripters to programmers, web development and open source.

We're thrilled about the reaction readers have been giving us about MacTech. In many people's minds, this is now the magazine that the insiders open up the second it hits their mailbox - and that's something that we're proud of.

MacTech has continued to refine and expand its production team, editorial staff and board - all are here to serve you. Let us know what you think. Drop me a note, or anyone on our team for that matter. While we may not respond to everything, we do read it all.

Enjoy the MacTech Sampler. If you've found just one or two articles that were really useful to you in this sample, imagine what a whole year will bring you. If you aren't currently a subscriber, subscribe today. It's RISK FREE. You can cancel at any time.

<http://store.mactech.com/sampler>

will get you a special offer.



Neil Ticktin  
Publisher  
MacTech Magazine



## ARTICLES & DEPARTMENTS

*Getting Started*

### **PHP, MySQL, And Forms**

by Dave Mark ..... 6

### **Exclusive File Access in Mac OS X**

by Paul T. Ammann ..... 15

*Mac in the Shell*

### **Screen**

#### ***Living In A Virtual World.***

by Edward Marczak ..... 22

*The Source Hound*

### **Knock Knock Knocking on LDAP's Door, PART 1**

by Dean Shavit ..... 30

*From the Source*

### **An Open Source Primer**

#### ***A Practical Guide to Using Open Source Software (OSS) on Mac OS X for the Non-Developer***

by Emmanuel Stein ..... 46

### **Deconstructing RSS 2.0**

#### ***Understanding How RSS Feeds Work***

by Dave Woolridge ..... 56

*Patch Panel*

### **Nagios on OS X, Part 1**

#### ***Installing and Setting Up One of the Best Net Monitoring Tools Around***

by John C. Welch ..... 68

## Communicate With Us

### Department E-Mails

#### Orders, Circulation, & Customer Service

cust\_service@mactech.com

#### Press Releases

press\_releases@mactech.com

#### Ad Sales

adsales@mactech.com

#### Editorial

editorial@mactech.com

(Authors only, no pr)

#### Accounting

accounting@mactech.com

#### Marketing

marketing@mactech.com

#### General

info@mactech.com

#### Web Site

http://www.mactech.com

.....

In this electronic age, the art of communication has become both easier and more complicated. Is it any surprise that we prefer **e-mail**?

If you have any questions, feel free to call us at 805/494-9797 or fax us at 805/494-9798.

If you would like a subscription or need customer service, feel free to contact MacTech Magazine Customer Service at 877-MACTECH

.....

We love to hear from you!  
Please feel free to contact us with any suggestions or questions at any time.

Write to letters@mactech.com or editorial@mactech.com as appropriate.

# MACTECH<sup>®</sup>

The Journal of Macintosh Technology

A publication of **XPLAIN** CORPORATION

## The Editorial Staff

**Publisher & Editor-in-Chief:** Neil Ticktin

**Associate Publisher :** David Sobsey

**Executive Editor:** Edward R. Marczak

**Editor-at-Large:** Dave Mark

**Business Editor:** Andrea Sniderman

**Editor-at-Large, Open Source:** Dean Shavit

**Managing Editor:** Dennis Bower

**Copy Editor:** Marianne Shilpa Jacobie

**Staff Writer:** Dharmendra Rai

## Xplain Corporation Senior Staff

**Chief Executive Officer:** Neil Ticktin

**President:** Andrea J. Sniderman

**Accounting:** Marcie Moriarty

**Customer Relations:** Susan Pomrantz

**Board of Advisors:** Steven Geller, Alan Carsrud

## Regular Columnists

**QuickTime ToolKit:** by Tim Monroe

**Patch Panel:** by John C. Welch

**The Source Hound:** by Dean Shavit

**Reviews/KoolTools:** by Michael R. Harvey

**AppleScript Essentials:** by Ben Waldie

**Mac In The Shell:** by Ed Marczak

## Board of Advisors

**Chairman:** Dave Mark,

Jordan Dea-Mattson, Steven Geller, Bruce Friedman, and Richard Kimes

## Contributing Editors

Michael Brian Bentley, Gordon Garb, Vicki Brown, Chris Kilbourn  
Marshall Clow, Rich Morin, Will Porter, Tom Djajadiningrat, Avi Rappoport,  
Andrew S. Downs, Cal Simone, Steve Sisak

**MacTech Magazine** (ISSN: 1067-8360 / USPS: 010-227) is published monthly by Xplain Corporation, 850-P Hampshire Road, Westlake Village, CA 91361-2800. Voice: 805/494-9797, FAX: 805/494-9798. Domestic subscription rates are \$47.00 per year. Canadian subscriptions are \$59.00 per year. All other international subscriptions are \$97.00 per year. Domestic source code disk subscriptions are \$77 per year. All international disk subscriptions are \$97.00 a year. Please remit in U.S. funds only. Periodical postage is paid at Thousand Oaks, CA and at additional mailing office.

**POSTMASTER:** Send address changes to **MacTech Magazine**, P.O. Box 5200, Westlake Village, CA 91359-5200.

All contents are Copyright 1984-2006 by Xplain Corporation. All rights reserved. MacTech and Developer Depot are registered trademarks of Xplain Corporation. RadGad, Useful Gifts and Gadgets, Xplain, DevDepot, Depot, The Depot, Depot Store, Video Depot, Movie Depot, Palm Depot, Game Depot, Flashlight Depot, Explain It, MacDev-1, THINK Reference, NetProfessional, NetProLive, JavaTech, WebTech, BeTech, LinuxTech, MacTech Central and the MacTutorMan are trademarks or service marks of Xplain Corporation. Sprocket is a registered trademark of eSprocket Corporation. Other trademarks and copyrights appearing in this printing or software remain the property of their respective holders.

## ARTICLES & DEPARTMENTS

### Distributed Computing

#### **Build Your Own Supercomputer From Your Macs Laying Around** ***How to Employ Tiger's Xgrid to Build Clusters and to Contribute to Grid/Cluster Projects***

by Mary Norbury-Glaser ..... 80

### AppleScript Essentials

#### **Providing Progress Feedback During Script Execution**

by Benjamin S. Waldie ..... 92

### Enhancing Applications Websites

#### **Adding Ajax to a Website**

#### ***Creating a Dynamic, User-friendly Website Interface is Simple and Straightforward***

by Andrew Turner ..... 102

### Focus Review

#### **CRYPTOCARD's CRYPTO-Server 6.3 for OS X:**

#### ***Eliminate Insecure Static Passwords***

by the MacTech Review Staff ..... 112

### Web Performance Testing

#### **Web Benchmarking 101**

#### ***A Guide to Stress Testing Your Website with ApacheBench and JMeter***

by Jin Lin and Emmanuel Stein ..... 116

### MICROSOFT | MAC IN THE ENTERPRISE

#### **Entourage 2004 Spotlight Support**

#### ***An IT Perspective: How Microsoft Entourage 2004 Now Takes Advantage of Spotlight***

by Brian Johnson and Andy Ruff ..... 122

# PHP, MySQL, AND FORMS

In our last PHP/MySQL exploration, we used the MySQL monitor to build a database, create a new table, then populate the table with data. We *did* use our PHP chops to query the table and spit out some HTML with the results, but that's a pretty bare-bones/hard-coded approach. A more common technique is to use an HTML form to populate the table. For example, suppose you were building an order-entry system for your sports memorabilia shop. You could use the MySQL monitor to enter all the data that represents your inventory. If you are really comfortable with the monitor, that's not a terrible approach. But, what if you want your assistant, who is a whiz at the web but has no MySQL experience, to manage the data? Creating a series of HTML-driven forms is definitely the way to go.

Before we dig into this month's sample code, I'd like to take a momentary detour to talk about a terrific new book I just got.

## Hacking Mac OS X Tiger

Back in the early days of Macintosh, there were just a very few Mac programming titles. One of the best of these books was called *Macintosh Programming Secrets*. Written by Scott Knaster and Keith Rollin, the book quickly became a cult classic. Well, Scott is back and his latest book is every bit as cool as *Macintosh Programming Secrets*. The long title is

*Hacking Mac OS X Tiger: Serious Hacks, Mods, and Customizations*. This is no idle boast. The book is chock full of wonderful insider info, and is just plain fun.

The book starts off with a collection of tips. For example, try holding down the option key, then pressing the volume up (or down) key on your keyboard. Cool! The Volume System Pref opens. There are tips for the Finder, the Dock, Dashboard, System Prefs, iTunes, and a lot more. To me, the book is worth the price of admission just for that part alone. But wait, there's more!

The second part of the book is called Mods, and explores more developer-oriented things. Things like Automator, Xcode, Property Lists, and Application Bundles. Lots of great info here, especially if you are relatively new to Mac development.

The third part of the book might be my favorite. It's called Hacks, and is full of, well, um, hacks! You'll customize dock icon behavior, hack some Dashboard

widgets, even redirect your web cam output to create a live video desktop.

All this stuff is presented in Scott Knaster's witty, irreverent style. I totally love this book! *Hacking Mac OS X Tiger* is part of Wiley's *ExtremeTech* series. You can find it at:

<http://www.wiley.com/WileyCDA/WileyTitle/productCd-076458345X.html>

And now back to our regularly scheduled programming...

## Implementing a Form with PHP

The key to implementing an HTML form is the `<form>` tag. Typically, the `<form>` tag will include both action and method attributes. The action attribute specifies the name of the file to which to send the form's results. The method attribute specifies how form data is passed along to the action recipient and is either GET or POST. The primary difference between GET and POST is that GET causes the form data to be embedded in the action's URL, while POST passes the data to the action as an input stream (via stdin, for example). GET forms can be bookmarked, since the data is embedded in the URL. POST forms cannot be bookmarked.

In general, I use POST as my default unless I specifically need a URL to be bookmarkable. You've undoubtedly clicked on your browser's reload button and been asked, "Do you really want to reload this form?" or something similar. The page in question is a POST request.

Good to know the difference between GET and POST, use whichever works best in any situation. Note that GET is the default. Note also that if your form contains passwords and the like, the GET URL will include the password data. Not a great idea if you will be passing the URL around.

Embedded between the `<form>` and `</form>` tags is the form content. The form content is made up of HTML formatted text, along with a variety of `<input>` tags. Each `<input>` tag represents an HTML form element, such as a popup menu, text field, non-echoing text field (for entering passwords), checkbox, radio button, submit button (a pushbutton whose purpose is to submit the form) or a reset button (a pushbutton used to clear the form).

An example will make this a bit clearer.

### A PHP Form Example

This example puts up a pair of text fields, with appropriate labels. One field is a standard text field, the second a non-echoing text field. We'll use the former to

enter your name and the latter to enter a password. We'll end the form with a submit button. When the submit button is pressed, the same page will be reloaded. The technique of taking two different actions on the same page of PHP code is a common PHP theme. While this is not a particularly useful example, it demonstrates this concept nicely.

Using BBEdit, or your favorite plain-text editor, create a new plain-text file and save it as *php\_form01.php*, then type the following source code into the file:

```
<html>
<head>
  <title>Sample PHP Form</title>
</head>
<body>
  <?php
    if ( !empty( $_POST[ 'name' ] ) ) {
      echo "Your name is {$_POST[ 'name' ]}.<br />";
      echo "Your password is {$_POST[ 'pwd' ]}.<br />";
    }
  ?>
  <form action="<?php $PHP_SELF; ?>" method="POST">
    Name:
    <input type="text" name="name" />
    <br />
    Password:
    <input type="password" name="pwd" />
    <br />
    <input type="submit" />
  </form>
</body>
</html>
```

Save the file and copy it into the *Sites* directory in your home directory. For example, I placed my copy of this code into */Users/davemark/Sites/*.

Once this is done, fire up your favorite browser and enter this URL:

[http://127.0.0.1/~davemark/php\\_form01.php](http://127.0.0.1/~davemark/php_form01.php)

Be sure to replace davemark with your own user name. Also, be sure to keep the tilda (~).

Once you enter the URL, you should see results similar to those shown in Figure 1. If you don't, be sure your PHP server is running. To do this, create a new, plain-text file named *test.php*, fill it with these 3 lines of code:

```
<?php
phpinfo()
?>
```

Drag the test file into your *Sites* directory and type this URL into your browser:

<http://127.0.0.1/~davemark/test.php>

Remember to replace davemark with your user name. If the large PHP table appears in your browser, PHP is running and you've got an error in your version of *php\_form01.php*. If the large PHP table does *not* appear, go back to the installation instructions for PHP and make sure your server is up and running.

# DON'T LET YOUR MAC WIND DOWN!

CLEAN UP YOUR MAC  
UNINSTALL APPLICATIONS  
BOOST PERFORMANCE  
REGAIN LOST DISK SPACE

INCLUDES FULL VERSIONS OF:  
TECHTOOL PLATINUM  
SPRING CLEANING  
DANTZ RETROSPECT EXPRESS



A Division of Smith Micro Software

**Available from your favorite retailers and catalogs. Learn more about CheckIt at [Allume.com](http://Allume.com)!**

CheckIt System Performance Suite is the most powerful tool to diagnose, clean, optimize, back up, and repair your Macintosh!

® TM, and © 2006 Allume Systems, Inc. Mac OS X 10.3 or higher is required.




Figure 1. The Sample PHP Form in action.

Click in the *Name* field and type your name. Next, tab over to or click in the *Password* field and type a password. Note that the password characters echo as dots. I'll show you why in a second.

Finally, click the Submit button. You should see something like the output shown in Figure 2. Notice that the name and password are both echoed and the form is then reloaded so you can try this again.




Figure 2. The name and password are echoed and the form is reloaded.

Let's take a look at the code, see how this works.

We start with the basic `<html>`, `<head>`, `<title>`, and `<body>` tags. Nothing new there.

```
<html>
<head>
  <title>Sample PHP Form</title>
</head>
<body>
```

The first thing we do in the HTML body is open a PHP tag and jump into an if statement. The if statement calls the PHP function *empty()* to see if the variable `$_POST['name']` is empty, that is, see if it has been set to a value.

Whenever you see a PHP function and don't know what it is, make your way over to <http://php.net>, type the function name in the *search for* text field, select *function list* from the *in the* popup menu, then click the arrow to do your search. The result of my search for the *empty* function is shown in Figure 3.




Figure 3. Searching for the empty function on php.net.

What the heck is `$_POST['name']`? What a weird name for a variable. The variable `$_POST` is an *associative array*. An associative array is an array that is indexed by name instead of by numerical index. For example, you might have an array of nicknames for all your friends, where `$nicknames['Fred']` is set to the string 'Stubby', while `$nicknames['Phil']` is set to 'Scooter'.

Note that the quotes are optional for single word strings. So you might see `$nicknames[Fred]`. Though that will work, leaving out the quotes will make your code a bit more cryptic, so use them unless you find yourself in a funky situation where they just get in the way (quoted string inside another quoted string, for example). And even then, be sure to comment your code so everyone can follow along.

`$_POST` is an associative array that contains an element for each of your form's input fields. So `$_POST['name']` contains the contents of the input element with the name "name". `$_POST['pwd']` contains the contents of the input element with the name "pwd". You get the idea.

The goal of the if statement is to see if the name field is empty. If not, we'll use echo to display the contents of the name field. Want to learn about echo? Go to php.net and type echo in the search field, search the function list. About 70% of the way down the page, you'll see a nice comment about using braces. This is worth reading. In a nutshell, you can use braces as separators when a space character just won't do. For example, in the code below, we used the braces to set off the name `$_POST['name']` without putting a space between the end of the variable and the period. This lets us put a period immediately after your name.

Note that we included the `<br />` HTML tag in our output. Remember, the PHP code spits out HTML as its output and then the HTML is passed back to the browser. The `<br />` tag is used to put a return after the name and again after the password.

```
<?php
if ( !empty( $_POST[ 'name' ] ) ) {
    echo "Your name is {$_POST[ 'name' ]}.<br />";
    echo "Your password is {$_POST[ 'pwd' ]}.<br />";
}
?>
```

# Now scheduling and contact management for your entire organization.



## Now | **Up to Date®** and **Contact® 5**

Calendar and Contact Software for Business, Groups, Families and You.



Is this project on schedule? When are you available to meet about the systems upgrade? Where are all the field techs today? When was the last time anyone talked to our biggest customer?

Virtually all groups live (or die) by their abilities to meet deadlines and keep track of their customers, prospects, and vendors. Few small companies or even departments of big companies have the tools they need.

Now Up-to-Date & Contact might just be the calendar and contact software for you. It's time-tested and used by more Mac-based companies than any other solution. And it's cross-platform-available for your PC users, too. It's easy to install and manage and simple for your employees to understand and use.

Using Now Up-to-Date & Contact you can schedule meetings for multiple users, view multiple, simultaneous calendars, and reserve rooms and resources. You can share contact information about your customers, prospects and vendors. And using our free server software you can set it up in minutes and share with users in the office or from anywhere with an internet connection.



Phone: 866-527-0556

Web: [www.nowsoftware.com](http://www.nowsoftware.com)

Call us now at 866-527-0556 or email us at [mactech@nowsoftware.com](mailto:mactech@nowsoftware.com) and we'll send you our free evaluation kit, including the book that will make it all easy, "Take Control of Now Up-to-Date & Contact" from Take Control books!

If the name field was filled, we'll already have printed the name and password before we get to this point. If the name field was empty, nothing will have been generated yet. Either way, we are now going to put up the form.

The `<form>` tag takes an action attribute and a method. We already know about the method. The action is the file you want served up with the results of your action. We could have created a second page that displayed the results of our form. But this self-referential approach is pretty common in PHP. We used the variable `$PHP_SELF` to have the form pass the results on to the same file, sort of like a function calling itself. Want to learn more about `$PHP_SELF`? It's not a function, so we'll need to search the whole site instead of just the function list. When you search all of php.net, you'll get a Google list of pages to look at. Here's a good place to start:

<http://us3.php.net/reserved.variables>

You'll need to search the page for `PHP_SELF`. Note that the `$` is not part of the variable name.

```
<form action="php $PHP_SELF; ?" method="POST">
    Name:
```

The two `<input>` tags specify a text field and a password field. Both do the same thing, but the password field echoes input as dots instead of the actual characters. Note that the name attributes of each `<input>` tag are what link the results to the `$_POST` associate array index.

```
<input type="text" name="name" />
<br />
Password:
<input type="password" name="pwd" />
<br />
```

Finally, the last bit of `<input>` for the form is the submit button.

```
<input type="submit" />
</form>
</body>
</html>
```

Type your name and password again and again. Each time you click submit, the page reloads, draws your name and password, and then redraws the blank form. Cool!

## Adding MySQL to the Picture

Our next step is to add MySQL to the picture. We'll build a database, using the MySQL monitor, create a new table, then use the form we created to populate the database. If you are new to the MySQL monitor, go back June's *Getting Started* column for a review.

Fire up Terminal and start the MySQL monitor. As a reminder, I set up an alias for mysql:

```
alias mysql='/usr/local/mysql/bin/mysql'
```

Here's my sequence for starting MySQL, building the database, and building the table:

```
Daves-Computer:~ davemark$ mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 16 to server version: 4.1.12-
standard
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the
buffer.
```

```
mysql> create database myaccounts;
Query OK, 1 row affected (0.40 sec)
```

```
mysql> use myaccounts;
Database changed
mysql> create table passwords(
-> name varchar(60),
-> pwd varchar(60),
-> userid int(10) auto_increment primary key );
Query OK, 0 rows affected (0.08 sec)
```

```
mysql> select * from passwords;
Empty set (0.09 sec)
```

In the above, I first created a database called *myaccounts*, set myaccounts as the current database, then created a table within myaccounts called *passwords*. Passwords contains a 60 character name, a 60 character password, and an auto-incrementing userid. We then verified that the passwords table was empty.

Let's switch back over to PHP and modify our code so we are writing our passwords into the table. Back in your plain-text editor, replace your existing code with this version:


```
<html>
<head>
<title>Sample PHP Form</title>
</head>
<body>
<p>Connecting to the database...</p>
<?php
    $host = 'localhost';
    $user = 'root';
    $pw = '';
    $db = 'myaccounts';

    $link = mysql_connect( $host, $user, $pw )
        or die( 'Could not connect: ' . mysql_error() );
    echo 'Connected successfully';
    mysql_select_db( $db );
    ?>
<p>If we got here, we've connected!</p>


<?php
    if ( !empty( $_POST[ 'name' ] ) ) {
        $tempname = $_POST[ 'name' ];
        $temppass = $_POST[ 'pwd' ];
        $query = "insert into passwords values (
'$tempname', '$temppass', 0)";
        $results = mysql_query( $query )
            or printf( "Query error: %s", mysql_error() );
    }
    mysql_close();
    ?>
<form action="php $PHP_SELF; ?" method="POST">
    Name:
    <input type="text" name="name" />
    <br />
    Password:
    <input type="password" name="pwd" />
    <br />
    <input type="submit" />
</form>
</body>
</html>
```

The first chunk of this code should be familiar from earlier columns. We use *mysql\_connect()* to connect to

# When a stranger calls.



Let Phlink take the call. Personalized greetings, voicemail, notifications, and much more.



Ovolab Phlink is the ultimate message center for your Mac. It **answers** phone calls and identifies callers using Caller ID and Apple's Address Book. It greets your friends with personalized messages. It **records** and stores messages on your computer – and even forwards voicemail to **email** as AAC audio attachments. Featuring multiple **voice mailboxes**, high-quality audio, Spotlight searching and fax capabilities, Ovolab Phlink makes

your telephone part of the digital hub! And you can fully customize Ovolab Phlink to do exactly what you need, using AppleScript: even set it up to call you back on your cell phone when important clients leave a message. Check it out now at [www.ovolab.com](http://www.ovolab.com).

**OVO**<sup>LAB</sup>  
software for the creative mind



the database and `mysql_select_db()` to make the database the current one for subsequent commands (just like `use myaccounts;` in the MySQL monitor).

```
<html>
<head>
  <title>Sample PHP Form</title>
</head>
<body>
  <p>Connecting to the database...</p>
  <?php
    $host = 'localhost';
    $user = 'root';
    $pw = '';
    $db = 'myaccounts';

    $link = mysql_connect( $host, $user, $pw )
      or die( 'Could not connect: ' . mysql_error() );
    echo 'Connected successfully';
    mysql_select_db( $db );
  ?>
  <p>If we got here, we've connected!</p>
```

The next chunk of code looks like what we did earlier in the column. Instead of using `echo` to print out the form fields, we build a query instead and use `mysql_query()` to fire it off. This will insert the row into the MySQL table. It is important to note that the `$query` line is a single line of code. It wraps in this column, but you should enter it as a single unbroken quoted string in your source code.

After the if statement, we close the database. Each time this page is loaded, the data base is reconnected, worked with, then closed again. Some people put the `mysql_close()` at the very bottom of the file, just to make sure they don't accidentally place code beneath it.

```
<?php
  if ( !empty( $_POST[ 'name' ] ) ) {
    $tempname = $_POST[ 'name' ];
    $temppass = $_POST[ 'pwd' ];
    $query = "insert into passwords values (
'$tempname', '$temppass', 0)";
    $results = mysql_query( $query )
      or printf( "Query error: %s", mysql_error() );
  }
  mysql_close();
?>
```

The rest of the code is pretty much the same as what you had above, the implementation of the form and the close body and html tags.

```
<form action="<?php $PHP_SELF; ?>" method="POST">
  Name:
  <input type="text" name="name" />
  <br />
  Password:
  <input type="password" name="pwd" />
  <br />
  <input type="submit" />
</form>
</body>
</html>
```

## Running the MySQL Form

Let's take this puppy for a spin. First, be sure you saved your source code and copy the file to your *Sites* folder. In your browser, type this link:

[http://127.0.0.1/~davemark/php\\_form02.php](http://127.0.0.1/~davemark/php_form02.php)

Note that I saved this new version of the source code as `php_form02.php`. Be sure to change *davemark* to your user name.

Once the table loads, enter a name and a password (see Figure 4) and click the **Submit** button. Note that each time the page reloads, the database is reconnected.




**Figure 4. Entering the data that will be written to the MySQL table.**

Once you've entered enough data, go back to Terminal and take a look at your table.

```
mysql> select * from passwords;
+-----+-----+-----+
| name          | pwd          | userid |
+-----+-----+-----+
| Franklin Roosevelt | Delano      | 1      |
| Dave Mark      | fudgiethewhale | 2      |
| Dark Helmet    | spaceballs   | 3      |
+-----+-----+-----+
3 rows in set (0.00 sec)
```

Cool! All the data is there.

## Until Next Month...

There is just a *ton* you can do with MySQL and PHP. If you don't have a PHP library already, I found the books *MySQL in a Nutshell* and *Programming PHP* from O'Reilly and *PHP 5 and MySQL* from Apress a big help in getting started.

Not sure what I'm going to do for next month's column. I've been reading the galleys to Mark Dalrymple and Scott Knaster's upcoming *Learn Objective-C* from Spiderworks. Hmmm, might be fun to explore a bit of that. Well, we'll see. For now, it's back to Lake Anna for a bit more summer laziness. See you next month. ☺



## About The Author



*Dave Mark is a long-time Mac developer and author and has written a number of books on Macintosh development. Dave has been writing for MacTech since its birth! Be sure to check out the new Learn C on the Macintosh, Mac OS X Edition at <http://www.spiderworks.com>.*

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!



# Exclusive File Access In Mac OS X

By Paul T. Ammann

Opening a file from classic Mac OS (pre Mac OS X) with `fsWrPerm`, `fsRdWrPerm`, or the default `fsCurPerm`, meant that any other application trying to open that same file with write access would not be able to do so. Usually, an `fsRdWrPerm` error would be returned when other attempts were made to open the file for write access, though attempts to open such a file for read-only access would succeed. This default behavior allows for one “write” and multiple “readers” of the file.

Mac OS X’s BSD subsystem does not enforce file read/write privileges in the same way as classic Mac OS. Opening a file for writing does not ensure other processes can not write to the same file. The default behavior of BSD allows for multiple “writers” to a single file. As a result, opening a file via `PBOpenDF`, `PBOpenRF`, `PBOpen`, `PBOpenFork`, `FSOpenFork`, `HOpen`, etc., on a local volume and passing in a permissions value of `fsCurPerm`, `fsWrPerm`, or `fsRdWrPerm` does not guarantee exclusive file access on Mac OS X. On Mac OS X, subsequent `Open` calls to open a file with write permission may succeed without error. Similarly, the `PBLockRange()` routines may not actually guarantee byte ranges that cannot be modified by other processes. Because these routines may return without error, you should check out the availability of exclusive file access

(see “Checking Availability of Exclusive File Access”) before making any assumptions about the underlying file access. If the “supports advisory locks” feature is not available, your application will not know if the file is already in use by another application.

AppleShare servers and Personal File Sharing on Mac OS X do enforce exclusive file access and range locking for volumes accessed over the network. However, this functionality is only available when accessing files over a networked file sharing connection and is not available to applications running on the server itself.

## Guidelines for Working with Non-exclusivity

You should realize that many applications relied on the behavior of the classic Mac OS File Manager to prevent multiple applications from writing to the same file (or to control write access through byte range locking). Since that behavior is not implemented in all versions of Mac OS X, some common workarounds that you may wish to use in your code are described below. BSD was designed without exclusive locks in order to prevent denial of service attacks in which one process opens a file with an exclusive lock which may be required by another process, effectively blocking the other process.



# MAXIMUM PERFORMANCE TAPE

## INTRODUCING MAXELL DLTtape® S4

Once again Maxell leads the field, becoming the world's first qualified manufacturer and supplier of DLTtape S4. Developed in partnership with Quantum, Maxell DLTtape S4 offers an 800GB native capacity (1.6TB compressed), and incorporates Quantum DLTSage™ WORM and Tape Security features for legal/regulatory compliance and data security. Additionally, Maxell NeoSMART™ technologies ensure ultimate recording reliability and data recovery, furthering Maxell's advance toward 10TB capacities and beyond. Maxell computer tape - the maximum performance data storage solution.



**maxell**  
[www.maxell.com](http://www.maxell.com)

## Checking Availability of Exclusive File Access

Mac OS X will enforce exclusive file access, i.e. one writer and many readers of a file, through its application frameworks, Carbon, Cocoa, and Java by enforcing BSD advisory locks as though they are exclusive. The “supports advisory locks” feature is defined if both the OS and file system for the volume in question support advisory locks. In this case, the default behavior of the application frameworks is to open files with exclusive access when opened as writable. Applications built on these frameworks automatically get this functionality and do not need to be modified. When the conditions are met to support exclusive file access, **PBLockRange** will also call down through to the BSD advisory locks. Since **PBLockRange** will be based on BSD advisory locks at this point, range locks can be applied to local files as well as those on file servers.

Since not all versions of Carbon on Mac OS X support exclusive file access nor do all file systems support BSD advisory locks, you should check a couple of things before making assumptions about the underlying file access behavior. You should only assume these features are available if the gestalt bit, **gestaltFSSupportsExclusiveLocks**, as well as the **GetVolParms** bit, **bSupportsExclusiveLocks**, are both set. For instance, the Carbon Framework File Manager routines support advisory locks by default when **SupportsExclusiveFileAccess** returns true.

```
#ifndef gestaltFSSupportsExclusiveLocks
#define    gestaltFSSupportsExclusiveLocks    15
#define    bSupportsExclusiveLocks            18
#endif

Boolean    SupportsExclusiveFileAccess( short vRefNum )
{
    OSErr    err;
    GetVolParmsInfoBuffer    volParmsBuffer;
    HParamBlockRec    hPB;
    long    response;
    Boolean    exclusiveAccess    =
false;

    err = Gestalt( gestaltSystemVersion, &response );
    if ( (err == noErr) && (response < 0x01000) )
    {
        err = Gestalt( gestaltMacOSCompatibilityBoxAttr,
&response );
        if ( (err != noErr)
|| ((response & (1 <<
gestaltMacOSCompatibilityBoxPresent)) == 0) )
            return( true );    //    Running on Mac
OS 9, not in Classic
    }

    err = Gestalt( gestaltFSAttr, &response );
    if ( (err == noErr)
&& (response & (1L <<
gestaltFSSupportsExclusiveLocks)) )
    {
        hPB.ioParam.ioVRefNum    = vRefNum;
        hPB.ioParam.ioNamePtr    = NULL;
        hPB.ioParam.ioBuffer    = (Ptr)
&volParmsBuffer;
        hPB.ioParam.ioReqCount    = sizeof( volParmsBuffer
);
        err = PBHGetVolParmsSync( &hPB );
        if ( err == noErr )
            exclusiveAccess    =
                (volParmsBuffer.vMExtendedAttributes
& (1L << bSupportsExclusiveLocks)) !=
```

```
0;
    }
    return( exclusiveAccess );
}
```

To check if a volume supports byte range locking via **PBLockRange** you should check the **bHasOpenDeny** bit returned from **GetVolParms**. See Technical Note FL37 ([http://developer.apple.com/technotes/fl/fl\\_37.html](http://developer.apple.com/technotes/fl/fl_37.html)) for more information about **PBLockRange** details.

## Common Workarounds

The following two techniques are frequently used to work around this issue on platforms that do not enforce exclusive file access.

### Lockfiles

A common approach used by many developers is to create a “lockfile” in the same directory as the file being opened. Whenever opening a file, “foo”, for instance, with write access you first try to create a lockfile, “foo.lock”, in the same location. If the file creation fails because the file already exists, you assume “foo” is already open by another application. Upon closing “foo” the application is also responsible for deleting “foo.lock”. A strength of this technique is it only makes one assumption about the underlying file system: the file creation operation is atomic. The obvious weakness is that since there is no OS support for this method, each application is responsible for implementing its own lock file mechanism, and there are no agreed upon standards or conventions for the naming of lock files.

### Edit a Copy

Another workaround relies on operating on a unique copy of the file. When a file is opened for editing, a duplicate of the file is created in **/tmp** directory with a unique name, and opened. When the user tries to save the document, the modification date of the original is matched against the date cached during the open of the file. If it has changed, you know the file was modified.

## MAC OS X Solutions

### BSD Advisory Locking

Although Mac OS X’s BSD subsystem does not implement provisions for exclusive write access, (i.e. mandatory locks), it does provide advisory locks. An advisory lock is a voluntary locking mechanism in which the underlying file system maintains a linked list of record locks. As long as your application and other applications respect the locks, only one application at a time will have write access to a particular file. Since these locks are voluntary it is the choice/responsibility of the application developer to respect or ignore advisory locks. If you would like to use advisory locks, this can be done by

following the instructions later in this article. By accessing files through the application frameworks (Carbon, Cocoa, Java), in versions of the OS supporting the advisory locks feature in frameworks, this will be provided automatically if you use the framework's file access methods.

Applications that call BSD file I/O functions directly will not gain this behavior for free, and therefore should be revised to set and respect advisory locks by specifying the appropriate flags when opening a file.

You should evaluate changing calls from:

```
fd = open( "./foo", O_RDWR );

to

fd = open( "./foo", O_RDWR + O_EXLOCK + O_NONBLOCK );
```

Where, **O\_EXLOCK** means Atomically obtain an exclusive lock, and **O\_NONBLOCK** means Do not block on open or for data to become available or Do not wait for the device or file to be ready or available.

## Implementing Advisory Locking

Anywhere you are calling the System.framework version of **open(2)** with write access, you should modify the parameters to include the "**O\_EXLOCK + O\_NONBLOCK**" flags, and handle errors being returned, where they may have succeeded in the past. The **open(2)** call will then fail if the file has already been opened for exclusive access by another process.

Advisory locks are associated with a process and a file. This has two implications:

- When a process terminates all its locks are released.
- Whenever a descriptor is closed, any locks on the file referenced by that descriptor are released.

## Implementing Byte Range Locking

BSD also provides advisory byte range locking support through the **fcntl()** function. By using advisory locking, your applications will be able to work in a cooperative with Carbon, Classic, and other applications in the future. In these circumstances, files should be opened with the **O\_EXLOCK** set and then ranges locked through the **fcntl()** call.

Stevens' "Advanced Programming in the UNIX Environment" (page 367) describes some techniques for using the UNIX service **fcntl()** to lock portion of a file for reading and writing (Stevens, 1999, p. 367).

Warning: A file lock request which is blocked can be interrupted by a signal. In this case the lock operation returns **EINTR**. Thus you may think you got a lock when you really did not. A solution is to block signals when locking. Another solution is to test the value returned by the lock operation and relock if the value is **EINTR**. Another solution, which we adopt here, is to do nothing about it.

**Recording Locking** is the term normally used to describe the ability of a process to prevent other processes from modifying a region of a file while the process is reading or modifying that

portion of the file. BSD provides access to its record locking mechanism through the **fcntl** function:

```
#include <sys/types.h>
#include <unistd.h>
#include <fcntl.h>

/*
 * Returns:
 *   -1 on error
 */
int fcntl(int filedes, int cmd, ... /* struct flock
*flockptr */ );
```

We'll start with the third argument (flockptr), which points to a flock structure:

```
struct flock {
    short l_type;          /* F_RDLCK (shared read lock), or
                           * F_WRLCK (shared write lock),
or
                           * F_UNLCK (unlocking a region)
                           */
    off_t l_start;         /* offset in bytes, relative to
l_whence */
    short l_whence;        /* SEEK_SET: file's offset is set
to
                           *
                           * l_start bytes from
beginning of file
                           *
                           * SEEK_CUR: file's offset is
set to its current
                           *
                           * value plus the
l_start (which can
                           *
                           * be + or -)
                           *
                           * SEEK_END: file's offset is
set to the size of
                           *
                           * the file plus the
l_start (which can
                           *
                           * be + or -)
                           */
    off_t l_len;           /* length of region, in bytes
                           * special case: if (l_len ==
0), it means that
                           * the lock extends to the
largest possible
                           *
                           * offset of the file. This
allows us to lock a
                           *
                           * region starting anywhere in
the file, up
                           *
                           * through and including any
data that is
                           *
                           * appended to the file
                           */
    pid_t l_pid;           /* returned when cmd = F_GETLK */
};
```

This structure describes:

- The type of lock desired (i.e. read lock, write lock, unlock)
- The starting byte offset of the region being locked or unlocked (l\_start and l\_whence)
- The size of the region (l\_len)

To lock an entire file, set **l\_start** and **l\_whence** to point to the beginning of the file (i.e. **l\_start= 0, l\_whence= SEEK\_SET**), and specify a length (**l\_len**) of 0.

Any number of processes can have a shared read lock on a given byte, but only one process can have an exclusive write lock on a given byte. To obtain a read lock the descriptor must be open for reading, and the region cannot have an exclusive write lock. To obtain a write lock the descriptor must be open for writing, and the region cannot have an exclusive write lock nor any read locks.

A young boy with light-colored hair, wearing a blue long-sleeved shirt and khaki pants, stands in a field of tall, golden-brown grass. He is looking up towards a dark blue night sky filled with stars. In the background, a single tree stands on the horizon under a soft, glowing light, possibly from a setting or rising sun or moon.

# Remember when the sky was the limit?

Get ready for a new way of thinking.



## **INTEL® SOFTWARE DEVELOPMENT PRODUCTS FOR MAC OS\***

Intel's commitment to Apple isn't just about hardware. It's also about offering development tools that use the power of the processor to take your applications to the next level. With support for Apple development environments and languages such as Xcode\* and Apple Frameworks, our products work with the tools you're already using today.

So whether you build applications for physics or financial analysis, Intel® Software Development Products give you the tools you need to transform power into performance.

**Download a trial version today.**

**[www.intel.com/software/products/nolimits](http://www.intel.com/software/products/nolimits)**

Copyright © Intel Corporation, 2006. All rights reserved. Intel, the Intel logo, are trademarks of Intel Corporation or its subsidiaries in the United States and other countries. \*Other names and brands may be claimed as the property of others.

Now, we will describe the second parameter (`cmd`) for `fcntl`. The possible commands and what they mean are described in the following:

- **F\_GETLK**: Determine if the lock described by `flockptr` is blocked by some other lock. If a lock exists that would prevent ours from being created, the information on that existing lock overwrites the information pointed to by `flockptr`. If no lock exists that would prevent ours from being created, the structure pointed to by `flockptr` is left unchanged except for the `l_type` member, which is set to `F_UNLCK`.
- **F\_SETLK**: Set the lock described by `flockptr`. If we are unable to obtain a lock (because of previous locks already granted for the region) then `fcntl` returns `-1` and `errno` is set to either `EACCES` or `EAGAIN`.
- **F\_SETLKW**: This command is a blocking version of **F\_SETLK** (the `W` in the command means “wait”). If the requested read lock or write lock cannot be granted because another process currently has some part of the requested region locked, the calling process is put to sleep. This sleep is interrupted if a signal is caught.

Be aware that testing for a lock with **F\_GETLK** and then trying to obtain that lock with **F\_SETLK** or **F\_SETLKW** is not an atomic operation. We have no guarantee that between the two `fcntl` calls some other process won't come in and obtain the same lock.

To save ourselves the trouble of allocating a `flock` structure and filling in all the elements each time, Stevens defines the function `lock_reg` and a number of macros that call it. Notice that the macros shorten the number of parameters by two, and save us from having to remember the `F_*` constants mentioned above.

```
#define read_lock(fd, offset, whence, len) \
    lock_reg (fd, F_SETLK, F_RDLCK, offset, \
    whence, len)

#define readw_lock(fd, offset, whence, len) \
    lock_reg (fd, F_SETLKW, F_RDLCK, offset, \
    whence, len)

#define write_lock(fd, offset, whence, len) \
    lock_reg (fd, F_SETLK, F_WRLCK, offset, \
    whence, len)

#define writew_lock(fd, offset, whence, len) \
    lock_reg (fd, F_SETLKW, F_WRLCK, offset, \
    whence, len)

#define un_lock(fd, offset, whence, len) \
    lock_reg (fd, F_SETLK, F_UNLCK, offset, \
    whence, len)
```

```
pid_t    lock_test(int, int , off_t , int , off_t );

#define    is_readlock(fd, offset, whence, len) \
    lock_test(fd, F_RDLCK, offset, whence, len)
#define    is_writelock(fd, offset, whence, len) \
    lock_test(fd, F_WRLCK, offset, whence, len)
```

```
int lock_reg(int fd, int cmd, int type, off_t offset,
int whence, off_t len)
{
    struct flock lock;

    lock.l_type = type;        /* F_RDLCK, F_WRLCK,
F_UNLCK */
    lock.l_start = offset;     /* byte offset, relative
to l_whence */
    lock.l_whence = whence;    /* SEEK_SET, SEEK_CUR,
SEEK_END */
    lock.l_len = len;          /* #bytes (0 means to
EOF) */

    return ( fcntl(fd, cmd, &lock) );
}
```

```
pid_t    lock_test(int fd, int type, off_t offset, int
whence, off_t len)
{
    struct flock lock;
    lock.l_type = type;        /* F_RDLCK or F_WRLCK */
    lock.l_start = offset;     /* byte offset relative to
l_whence */
    lock.l_whence = whence;    /* SEEK_SET, SEEK_CUR,
SEEK_END */
    lock.l_len = len;          /* #bytes (0 means to EOF) */
    if (fcntl(fd, F_GETLK, &lock) < 0){
        perror("fcntl"); exit(1);}
    if (lock.l_type == F_UNLCK)
        return (0);           /* false, region is not locked
by another process */
    return (lock.l_pid);       /* true, return pid of lock owner
*/
}
```

There are three important rules regarding automatic inheritance and release of record locks:

- Locks are associated with a process and a file. When a process terminates, all its locks are released. Whenever a descriptor is closed, any locks on the file referenced by that descriptor for that process are released.
- Locks are never inherited by the child across a fork (otherwise we could end up with two processes sharing a write lock)
- Locks may be inherited by a new program across an `exec`. This is not required by BSD and is therefore machine dependent

## References

Stevens, Richard W. (1999). *Advanced Programming in the UNIX Environment*  
Massachusetts: Addison Wesley Longman, Inc.  
ISBN: 0201563177



## About The Author

*Paul Ammann has been working in IT for almost 20 years. He is happily married to his wife Eve for 6 years. He finds writing the author's bio the toughest part the article.*

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# SCREEN

## LIVING IN A VIRTUAL WORLD.

**H**ello ladies and gentlemen! Step right up! Don't be afraid. Gather 'round! Come closer...that's it. Did you know that there are those who are hooked on the CLI? That some of us have to work on a monitor with less than 1024x768 of space? Or perhaps you know people that use ssh to perform a lot of work remotely? Today, I'll demonstrate the magic of "screen". For many, many reasons, screen is an indispensable utility. It's the answer – to several questions, for anyone doing command-line work, from the 12" iBook user, to the dual-G5 30" monitor owner. Watch, and be amazed!

### The more things change...

My first job after graduation was at Computer Associates ("Hello, Doug!...who wouldn't know a Mac if he mistakenly walked into an Apple store!"). While CA has attempted to make a push into the PC world, in their heyday they had some of the best selling software for IBM "big iron" - mainframes. Of course, after thinking that I was going to get a job programming games or graphics (not that I had any real experience to speak of), this was an awakening into the things that companies really need. So, I sat down at my desk, and took a gander at the dumb-terminal green screen that was sitting *on* my desk, staring back at me. Good 'ol 5250 in all its glory. After a little introduction to the world of REXX, JCL and HLLAPI, I thought, "it would be great if I had another terminal here, so I could compare code on one screen without having to quit out and get back into what

I'm working on." Well, Mr. Doug, who *was* a mainframe guy, showed me a great trick shortly after.

"Hey, Ed, I want to show you something really cool. Look at my terminal. Here, I'm logged into System B, but if I press Fn-F2, there's e-mail, and Fn-F3 brings up my text editor. And when I press Fn-F1, there's System B again! We use a system here called *virtual terminals*. Did anyone ever show you that?" Well, no, they hadn't. Now I had a way to be logged into several systems at once and not lose my place when I needed to switch to another. All with a single physical terminal sitting on my desk.

### The more they stay the same!

What does this tale of mainframe tutelage have to do with modern day OS X? Well, virtual terminals, of course! Thanks to the hard work of the GNU volunteers that have contributed to "screen", we have fantastic virtual terminal facilities on our Macs (and Linux, and IRIX, etc.). So what, right? I can open up multiple terminal *windows*! This is 2005, buddy, and I'm not working on some crummy old green-screen. I've got a dual-

mega-hyper-super-secret-quad-X-on-Intel-benchmark-smoking-egg-frying-electricity-sucking machine! I think 16GB of RAM can handle a terminal window or two!

Whoa! I'll show you why screen is more than just a way to have multiple terminals. Apple sees the light on this one, too: screen has shipped with OS X since 10.2. Thanks to this, we can even skip any discussion about acquiring the source or a package and getting it running on your system. We can get right to the good stuff!

## What's that in the road? A Head?

What exactly is a virtual terminal? Instead of having physical screens attached to a system for each session, screen allows us to create a session that lives virtually. That is, not attached to any physical console. To even drive that point home, screen allows us to detach the session, and pick it up from anywhere else! Even better, you can have multiple people attached to a single session from anywhere, like a poor man's VNC.

Unix already supports job control. However, once you background a job and logout, you can't bring it back to the foreground. More than that, if you don't take special measures, a background job will die with the parent shell when you log out. No longer will you have to say, "I'm Mike Jones and I use nohup." (Especially if your name isn't Mike Jones).

For the impatient, pop open Terminal.app (or, the excellent iTerm, GLTerm, or other favorite), and type `screen`. You should get a simple welcome screen as shown here:

```
Screen version 4.00.02 (FAU) 5-Dec-03
```

```
Copyright (c) 1993-2002 Juergen Weigert, Michael Schroeder  
Copyright (c) 1987 Oliver Laumann
```

```
This program is free software; you can redistribute it and/or  
modify it under the terms of the GNU General  
Public License as published by the Free Software Foundation;  
either version 2, or (at your option) any later  
version.
```

```
This program is distributed in the hope that it will be  
useful, but WITHOUT ANY WARRANTY; without even the  
implied warranty of MERCHANTABILITY or FITNESS FOR A  
PARTICULAR PURPOSE. See the GNU General Public License for  
more details.
```

```
You should have received a copy of the GNU General Public  
License along with this program (see the file  
COPYING); if not, write to the Free Software Foundation,  
Inc., 59 Temple Place - Suite 330, Boston, MA  
02111-1307, USA.
```

```
Send bugreports, fixes, enhancements, t-shirts, money, beer &  
pizza to screen@uni-erlangen.de
```

```
[Press Space or Return to end.]
```

Once you press space or return, you'll get back a prompt. That's it, you're running under screen! If you expected something fancier, sorry to disappoint, but you'll see that it's worth it!

`^a` (ctrl-a) becomes the important key when using screen. You use it to give screen commands. Now that

you're in screen, here's a quick demonstration. Type `^a`, followed by `:"` (that's `^a`, let go, type the colon). You will see the command mode prompt, ready for input – much like vi. Type `"caption always"` and press return. You'll see a status line along the bottom of your terminal. It should read, in reverse video, `"0 bash"`. `"0"` is the window number and `"bash"` is the default window name. Let's give it something more descriptive. Type `^a-A`, and you'll be prompted with:

```
Set window's title to: bash
```

Backspace over `"bash"` and fill in your own title for this window – I'm going to call mine `"window0"` – and press return. While you're there, get a file listing with `ls -l`, just to fill the window. Now for the magic: type `^a-c`. This will create a new window. Your file listing will seem to disappear, and the status line will change to read `"1 bash"`. This is simply a new terminal – all in the same terminal window. You can vaguely compare this to fast user switching without the cube effect, or switching users. `window0` is still there, of course. Let's give this window a title right now. `^a-A`, and I title mine `window1`. On this window, get a process listing with `ps ax`. Great. Now, how about another window? `^a-c`, and you've got `"2 bash"`. I'm going to title this one `"topwin"`. Of course, then I'm going to run top:

```
top -ocpu.
```

It's great that we've created a few windows, but it's even more useful if we can move between them! There's a few ways to do that. Here's a quick list, followed by an explanation:

`^a-n` – next window

`^a-p` – previous window

`^a-#` – Jump to window number `"#"`

`^a-^a` – Jump to last window (like a 'back' button)

You can step through each window in order with `^a-n`. Try it: you'll wrap to `window0` and see the file listing that you left there. Press `^a-n` again and you'll see `window1` with its old process listing. One more time and you'll get our top display running on `topwin` (you left it running, didn't you?).

One more trick in the magic show. Press `^a-d`. This will **detach** you from this session of screen. You'll be back to where you were before starting screen, with the exception of a note that says, `"[detached]"`. Now log out. Yep, you heard me: Apple-Shift-Q and log out. This normally crushes all processes for your user. But screen will protect us here. Now, log back in. Open a terminal, and type `screen -r`. The `-r` flag instructs screen to "re-attach" to a detached screen, rather than start a new session. When you press return, you'll be magically transported right back to where you left off. Yeah, it's magic.

From Macs to mobiles, devices to Nintendo DS™ -  
wherever your site goes, Opera can too.



Follow the standards.  
Break the rules.  
Test in Opera first.

Choose Opera for better  
accessibility and standards  
support.

Download the **free**  
**Opera browser** -  
[www.opera.com](http://www.opera.com)



# No! It's "What's that in the road ahead?"!

Anxious to learn more, are you? It would be difficult to present every single thing that screen can do in one column. Even two columns would be tough. So here are my favorite, practical tricks of screen, and how to perform them.

1. Multiplex Terminal.app – Ah, the wonders of modern technology! Look at this miniature marvel, the 12" PowerBook! Light and portable, it's everything you need in a traveling laptop. That is, of course, besides the screen real estate. One Terminal window looks great, but how many are you going to open? Just one, of course! Then run `screen`, and `^a-c` to your heart's content. Even on a larger monitor, how many individual Terminal.app windows do you want?
2. Go home – Nothing like having a personal transporter to get you from place to place. Well, aside from the fact that *they don't exist yet!* But there is a way to end your workday on time and continue later on. I had a large copy job to run at the end of the day: over a Terabyte. But I was at the client's site, and not at home. Everyone else had gone for the day. Here's the trick: use a machine at the client's site that is accessible via ssh. Pop open Terminal.app and immediately run `screen`. You can then start your copy job, and disconnect using `^a-d`, and go home. Once you're home (or wherever you're going), ssh in and run `screen -r` to reconnect and check on, or finish, the job. Why sit and stare at file copy status when you could be using that time to travel home?
3. Protect against an unexpected disconnect – In a world of broadband and high-speed access, there are still some secrets. The secret of the unreliable ISP is the most insidious. It's true! You may ssh into a remote machine over the Internet, only to find that you unceremoniously get dropped - Connection reset by peer, connection closed. Dial-up still exists, my friends! If you're managing via dialup, or find that you get dropped connections, run `screen` immediately after ssh-ing into a remote host. This way, if you get dropped while you're in the middle of editing an rc file with `vi`, you simply ssh back in, and type `screen -r`. You'll be right where you left off.
4. Have screen setup my environment for me – Amaze at the automatic setup! Swoon over the customized messages! `screen` looks for a `.screenrc` file in your home directory. You can customize the way screen starts up. Here's my startup file:

```
defscrollback                                20000
activity 'Activity in local window %n'
defmonitor                                    on
nethack                                        on
startup_message                               off
caption always "%{= g}%50-%t%{= g} %{= r}| %{= y}%c: %s %Y-%m-%d
%{= r}| %{= g}%W"
screen                                         irc
screen                                         -t                                     rootwin
screen                                         -t                                     remotewin1
screen                                         -t                                     remotewin2
screen                                         -t                                     shamewin
screen                                         -t                                     localwin
screen                                         -t                                     sstun
```

Here's the explanations:

**defscrollback** defines the scrollbar buffer for each window.

**activity...** defines the activity message. This goes hand-in-hand with the following line, `defmonitor on`. With this set, when a window you're not watching has activity in it – like when you're waiting for a compile to complete...or your irc session to show some activity – screen will display a status message to let you know. You can customize the message by following the directive with a string. `"%n"` will substitute the window number.

**nethack on** is in my `.screenrc` file because, well, I'm a geek! With this option on, the standard screen messages are replaced with nethack-like ones. If you've just started using screen, this has the potential to confuse you. Otherwise, it's fun!...if you're me, I suppose.

**startup message off** simply suppresses the intro screen.

**caption always**, which we saw a simple example of earlier, can also accept a string that will define the status-line display.

The final **screen -t...** lines simply fire up new windows for me, titled with the title I supply using the `-t` switch.

There's much more about all of these features in the screen man page.

5. Screen sharing – There are many times that people call me for help, and solving the problem involves the command-line. Well, I really want them to watch what I'm doing so they can learn. For this exercise, I may be in place A, they're in place B and the server is in place C. Once again, I sprinkle a little magic dust over my keyboard and I can type in the same session as my remote friend! Of course, I'm just using screen's `-x` flag. Much like the `-r` switch, `-x` also re-connects to a session, although one that is not detached. It's called multi-display mode. Here's the easy way to use it. Have person A ssh into the server and start screen. Have person B ssh into the same server as the same user and run `screen -x`. Person A and person B will now be sharing the same session. Either user can type and both will see all screen output. What's nice about this is that if there are any sensitive passwords to type, and only one party knows them, you can let the prompt come up and let the appropriate person type it.

## All good things...

Those 5 uses are the most common ways I use screen. Of course, there's plenty more in its top hat of tricks.

Since you can detach from a running session of screen, what happens if you start a new session of screen after detaching? Well, you'll have two separate sessions of screen running. If you detach

from the second session, and try to reconnect, you'll realize that you have to tell screen *which one* to reconnect to:

```
$ screen -r
There are several suitable screens on:
  3848.ttyp1.Jack-Kerouak (Detached)
  574.ttyp2.Jack-Kerouak (Detached)
Type "screen [-d] -r [pid.]tty.host" to resume one of them.
```

If I want to reconnect to the second session, I can tell screen that with the command `screen -r 574`. You can also find out this list anytime by using the `-list` flag:

```
$ screen -list
There are screens on:
  3848.ttyp1.Jack-Kerouak (Detached)
  574.ttyp2.Jack-Kerouak (Detached)
2 Sockets in /tmp/usccreens/S-marczak.
```

Note that the sessions here are exclusive to a single user id. You won't pick up the detached sessions of other users. If something goes awry, and instead of "Detached" one session says "Dead", you can clear it out with the `"-wipe"` flag: `screen -wipe 3848`.

screen will let you split your terminal. Create a few screens (`^a-c`), create some output on each, and then press `^a-S`: Boom! Split-screen! `^a-Tab` hops between segments. Each `^a-n` (or `-p`, or `-#`) is independent of the other window. So, you could have top running in one split, with a long running app in the other. Of course, on my local Mac, I'd probably just open another terminal window, but this is pretty handy on a remote system. `^a-X` removes the region that currently has the focus. This is why I never use Terminal.app's split bar. Thank you, screen!

You can screen into an app immediately by passing the program name into screen. `screen vi somefile.txt` will start vi under screen. This avoids spawning another copy of your shell, too. Naturally, all of the meta keys are then active, and you can create another window if you do need some other functionality.

screen can perform logging for you with the `^a-H` keystroke combo. Again, this is of more use on a remote system that you may detach from, but want the output of some long running command. One pressed, `^a-H` starts logging, tells you the name of the file it's logging into, and you're done. Pressed again turns off the log. Thought you'd like to know.

screen also sports a history buffer with copy and paste. Nethack mode gives some particularly good messages here. I'll be honest, this is a feature I barely use anymore locally. I have Terminal.app's unlimited scrollback buffer and a mouse. But remember, you can run screen and then disconnect. All of the output will accumulate in screen's scrollback, not on your terminal while in this state. `^a-]` gets you into copy/scrollback mode. Use the arrow keys to move around. Additionally, `^u` and `^d` will page up and down respectively. The escape key will bail you out.

There are plenty of other options and expansions on what is contained in this month's column. The man page happens to be excellent, so once you're comfortable with the basics, go dig in! One thing that is *not* possible as-is with the Apple-supplied screen is the ability to do anything fancy with multi-user. ACLs only work if the screen binary is set-uid.



Connect • Communicate • Collaborate • Securely

**Because meetings are more productive when people are on the same page... of their calendars**

**Get the message**



**Kerio MailServer**  
A groupware alternative to Exchange that syncs calendars, contacts and email with Entourage and Outlook. Integrated anti-spam and McAfee virus filtering provide secure, junk-free email for users on any platform.  
Mac OS X • Linux • Windows

[www.kerio.com](http://www.kerio.com)



© 2006 Kerio Technologies, Inc. All rights reserved. All other trademarks are property of their respective owners.

# Speed... Strength... Solutions!

**Portable and Desktop External Storage Solutions**  
**Up to 1.5 Terabytes**  
**Starting at just \$99.95**

**Other World Computing**

*Providing Solutions for Today's Technology Users*

**www.MacSales.com**

**Fastest  
Connection  
Speeds!**

Serial-ATA  
FireWire 800  
FireWire 400  
USB



## **Mercury™ On-The-Go**

FW800/400, USB 2.0,  
or SATA. The First 160GB  
2.5-inch bootable, shock  
resistant, fully bus powered  
storage solution on the market Perfectly designed  
for the mobile work force: Data, Digital Photos, Music



The **ONLY**  
Single Drive  
Solution up  
to 750GB!

## **Mercury™ Elite-AL Pro**

World's best aluminum enclosed  
storage solution. Single Drive  
models up to 750GB, Dual-  
Drive RAID's up to 1.5 Terrabytes  
(1500GB). Available SATA, FireWire  
800, FireWire 400, USB 2.0  
interface options. High-  
Performance & A/V Certified.  
Industry Leading 2 Year Warranty.



## **Neptune™ Value Done Right**

Super quiet fanless FireWire design  
with data transfer over 40 Megs per  
Second. Quality, reliable FireWire  
storage up to 500GB at an ultra  
economical price.



## **Mercury™ Elite Pro Classic**

Award winning High Speed, Super  
Silent bootable backup and external  
storage in state-of-the-art shell up to  
500GB. FireWire/USB Engineered for  
Audio, Visual, General Data, Graphics,  
and Backup



**Other World Computing**

*Serving the Computer Universe since 1988*

Visit [macsales.com](http://macsales.com) 800.275.4576

**Own the Future TODAY!**

Select your solution today and save  
[www.macsales.com/firewire](http://www.macsales.com/firewire)



Naturally, any set-uid binary brings security implications, so I like the default of 'off' for this. If multiuser support with fine-grain ACLs is what you need, get out of screen, get root, and add the suid bit to screen: `chmod 4755 /usr/bin/screen`. Run screen again, and get a screen command line with ^a: and type multiuser on. Then, grant another user full rights with `acladd username`. Then, "username" from the same system can access your session with "`screen -x yourusername/`". Otherwise, you can always set up a dummy account that both users can access, let on run screen first, and the second can connect with `screen -x`. Hit the man page for more information regarding ACLs.

Is there more wonderment? Isn't there always?!? While Apple does place the source to screen on the Darwin source site, there really isn't too much they had to change. The original source compiles cleanly and runs nicely. And just when I was about to say that there was a dearth of documentation regarding screen, I find these while Googling about: <http://computing.ee.ethz.ch/sepp/screen-3.9.9-to/> and <http://rucus.ru.ac.za/docs/screen/>. The former is a deeper look at most commands, while the latter is a really nice way to ease yourself into screen.

## Must come to an end!

Is screen actually magic? I hope it feels like it to you. I think all technology seems like a small bit of reality (physics, electrical engineering, etc.) and magic. The people who created and continue to work on screen certainly have my humble thanks (plus, I sent them some Nastro Azzuro when I first found screen). The real magic of screen is certainly in the code itself; it allows me to be in more than one place at a time!

Next month, we'll get back to basics. bash basics: scripting for the admin.

MI

### About The Author



*Ed Marczak, owns and operates Radiotope, a technology consulting company that implements mail servers and mail automation. When not typing furiously, he spends time with his wife and two daughters. Get your mail on at <http://www.radiotope.com>*

**There are times when you have to take your desktop with you...**



iLUGGER Bags  
for ALL G5 and Intel iMacs

### Pro Series Bags



G5  
Tower  
Lugger



23"  
Cinema  
Lugger

**NEW!**



30"  
Cinema  
Lugger

If you need uncompromising performance **To Go**  
Call us or visit our website

**LTA PROJECTS • [www.ilugger.com](http://www.ilugger.com) • 877-897-5158 ext 2**

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# **KNOCK KNOCK KNOCKING ON LDAP's DOOR, PART 1**

If there is such a thing as a tradition in the relatively new IT profession, the widely held belief that Apple Server solutions don't have an adequate, scalable, or enterprise-worthy implementation of Directory Services is about as close as it gets. Directory Services are generally methods of organizing and aggregating user, group, computer, and resource information, into a logical and accessible hierarchy, along with a strategy of maintaining and enforcing access control over those resources. Such traditions are somewhat related to assumptions and their evil cousins known as hysterias which are similar to beliefs that Netware Sysops often held about the superiority of eDirectory (Novell's mature and spectacular Directory Service) over Active Directory when it was released in 2000. Those Novell Sysops are probably not quite as snooty now that Active Directory dominates Enterprise IT.

## **Inside the Black Box**

While Novell was polishing its Netware Directory Services into what is now called eDirectory, and Microsoft was busy readying Active Directory, Apple was still mired in the throes of what I can only fairly call "faking it." AppleShare IP 6 and OS X Server 1.2 combined with Macintosh Manager for OS 9 and earlier, superimposed a user and workgroup authorization model on top of a simple database that ran on a workstation OS (AppleShare IP), or synchronized up alongside of directory

service far too advanced for it (NetInfo in OS X Server 1.2), along with policies for access control to local and network resources that worked well in a school environment for several reasons, one of which was the cartoonish appearance of the windows and icons. Macintosh Manager has now gone the way of other soon-to-be extinct technologies with the release of Tiger Server (it's not included, unless you're upgrading from a Panther Server installation).

Starting with Panther Server, Apple moved into the LDAP age of open standards and Open-Source software with Open Directory version 2, and its successor in Tiger Server, Open Directory version 3. LDAP, which is an acronym for

Now with Full Support  
for Universal Applications

3,248 hours typing code  
184 hours finding that one bug  
142 hours of meetings  
108 pizzas  
14 cancelled weekend trips  
11 all-nighters  
**1 call protects it all!**

**HASP**  
SOFTWARE DRM


**#1 Choice of Software Developers**

The new HASP family of products is the next generation in protection ensuring the highest level of security for your software. It provides an easy set of tools to automatically protect your software and implement new and innovative licensing options.



**Strongest anti-piracy solution:**

Our powerful, 128-bit AES encryption provides a strong locking mechanism that ensures you get paid for every copy of your software.



**Best overall solution as rated by an independent testing lab:**

In KeyLabs tests, HASP HL outperformed the competition in security, ease-of-use, flexibility of tools and functionality, and compatibility across platforms and environments. A copy of the full KeyLabs report is available from our Web site at [www.Aladdin.com/MacTech](http://www.Aladdin.com/MacTech).

**Robust envelope:**

Automatically protects Windows executables, DLLs, .Net and Mac applications, securing the weakest link and ensuring your valuable IP remains unreachable. Multiple protection layers provide unequalled security.



| Category                                              | Weight | Aladdin<br>HASP HL | Sentinel<br>UltraPro |
|-------------------------------------------------------|--------|--------------------|----------------------|
| Security                                              | 35%    | 85.1%              | 60.9%                |
| Ease-of-use                                           | 25%    | 93.7%              | 77%                  |
| Licensing                                             | 20%    | 75%                | 75%                  |
| Flexibility of tools<br>and functionality             | 10%    | 93.3%              | 71.2%                |
| Compatibility across<br>platforms and<br>environments | 10%    | 93.3%              | 71%                  |
| <b>Overall Evaluation</b>                             |        | <b>87.3%</b>       | <b>69.8%</b>         |



Request a Software Developer Kit today at [www.Aladdin.com/MacTech](http://www.Aladdin.com/MacTech)

**Aladdin**  
SECURING THE GLOBAL VILLAGE

North America: 1-800-562-2543, 847-818-3800 • UK • Germany • Israel • Benelux • France • Spain • Asia Pacific • Japan

© 2005 Aladdin Knowledge Systems, Ltd. All rights reserved. Aladdin and HASP are registered trademarks of Aladdin Knowledge Systems, Ltd. Windows®, Mac OS®, Linux® are trademarks or registered trademarks of their respective holders.

Lightweight Directory Access Protocol, has, at the very least, become an agreed-upon standard that eDirectory and Active Directory aspire to, or emulate, or at least talk to. Apple's Open Directory, however, doesn't have to aspire, because it is LDAP, OpenLDAP (www.openldap.org) to be precise. Tiger Server's LDAP implementation is based on OpenLDAP 2.2.19 and Berkeley DB version 4.2.52 (which is used to actually store the directory data).

Despite the fact that many of the initial-like components of LDAP are familiar buzz words in the ears of IT professionals "DN" for *distinguished name*, "CN" for *common name*, "O" for *organization name* and "OU" for *organizational unit name*, and "DC" for *domain component*, LDAP is still often considered a black box, something that's addressed, mapped to, queried, searched or occasionally edited or backed up. Aside from Novell's Console One eDirectory administration application, which has a general hierarchical view of the Directory contents, very few Directory Service tools, including Apple's own Workgroup Manager, do much to shine a light on the contents of the black box, or how it's organized. Luckily, there's an Open-Source tool that allows the curious admin a view of Apple's Open Directory with a visual arrangement that follows the conceptual structure.

## Set up the LDAP Server

Before you can use tools like PHPLDAPadmin to start examining the structure of an LDAP Open Directory Master, you need to make sure that your installation of Tiger server is functioning in the "Open Directory Master" role. You can easily examine the Open Directory data in a local NetInfo domain using the NetInfo Manager utility in /Applications/Utilities. Open up Server Admin, and click on the "Open Directory" service section on the left. You should see the following:



**Figure 1. Open Directory Overview in Server Admin**

If the overview indicates that your server's "Connected to a Directory System," then you should change the role to "Standalone Server," reboot for the sake of superstition, then

change the role to "Open Directory Master." If your server is a "Standalone Server," then you're good to change the role to Open Directory Master as well. Just like an ordinary Tiger workstation, Tiger Server, when configured as a Standalone Server, simply stores its Directory data in its local NetInfo (nldb) database. It isn't until Tiger Server is "promoted," to borrow a Microsoft Server term, that the OpenLDAP Server daemon (slapd) starts.

There is, however, one big *gotcha* to be aware of: in order for the LDAP configuration to finish properly, the DNS service needs to be set up to a "T" so that the server can resolve its own host name and reverse its IP address without errors, otherwise LDAP may not behave as expected, nor will the Kerberos KDC (Key Distribution Center) process start up automatically as it should. Luckily, setting up the DNS service properly has never been simpler than with Tiger Server; you may have heard that some sysadmins think that the new DNS GUI in Tiger Server is a step backward from the one in Panther, and while I agree, Apple certainly knows the needs of its users better than little 'ol me, but somehow I keep wishing for an "expert" (look at the wording here. Needs something added for it to make sense.) DNS editing mode that would let me easily do thing like edit reverse lookup records without having to go back into the terminal to configure DNS, as I did in the days of Jaguar Server.

First, open Server Admin and click on the DNS service on the left. In the window on the right, click the "Settings" tab at the bottom. Then, under the general tab on the top, go ahead and uncheck the "Zone Transfers" box (which is used only when there's a slave DNS server working with the server's DNS) for extra security, so that your DNS server can't be a victim of "spoofing" whereby a malicious process attempts to exploit the DNS replication mechanism.




**Figure 2. Disabling "Zone Transfers" for the DNS Service**

Second, click on the "Zones" tab and enter the name of your domain. Notice that my default, the DNS service GUI will repurpose the computer/Bonjour name of your server as the hostname. If at this point you want the hostname to differ, it's your opportunity to change it. And that's really about all there is to it; Apple has made it easy. Now, just make sure that the IP address of your server's set as the first DNS server in the list in System Preferences > Network > Your Ethernet Interface.



## WITHOUT EMAIL, FTP OR SLOW UPLOADS.

Whether it's in pro photography, print and prepress, or at the country's leading creative agencies, SeeFile 2.0 software for OSX is revolutionizing how media files are shared and organized. SeeFile uses the Web instead of client software to make uploading, downloading and reviewing jobs a snap.

SeeFile turns a Mac's file system into a continuously-updated Website with complete user control. Accessible from anywhere via Safari, Firefox or Windows IE, folders look like folders, automatically generated thumbnails can be selected to get a high-res preview of each image, and our integrated messaging allows online discussion of specific files or whole jobs.

Private logins allow separate user areas as well as multiple levels of responsibility and access. Easy logo branding also means that you can set up custom-looking sites for new clients in seconds.

SeeFile is priced at a fraction of traditional workgroup digital asset managers, and runs on any G4, G5 or Intel OSX system. OSX Server is not required - the software takes advantage of the Apache server built into every OSX system. SeeFile's extensible scripting and hotfolders, as well as PHP, MySQL and XMP metadata support, make it easy to integrate with third party workflows; it's already pretty powerful right out of the box.

We're looking for more good Mac consultants to partner on installations, grow their business with us and help spread the word. We're easy to work with, and provide old-style phone contact as well as in-person and VNC tech support. Contact us today and see what's possible.

[www.seefile.com](http://www.seefile.com) | 617-262-2464 | [info@seefile.com](mailto:info@seefile.com)

**see**file™  
LIVE FILE SHARING •  
ONLINE APPROVALS




Figure 3. Zone Setting Screen.

At this point everything should work if you save your changes and then click the big green “Start Service” button, but I’ve developed a ritual that dates back to the OS X 10.1 days. It’s called “digging DNS,” and consists of two simple tests to verify that the DNS is humming along. First, test the forward lookup record like so:

```
mostsvr:~ mostadmin$ dig mostsvr.macworkshops.com

; <<>> DiG 9.2.2 <<>> mostsvr.macworkshops.com
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id:
39577
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY:
1, ADDITIONAL: 0
;; QUESTION SECTION:
mostsvr.macworkshops.com.      IN      A
;; ANSWER SECTION:
mostsvr.macworkshops.com.  86400 IN      A
192.168.0.85
;; AUTHORITY SECTION:
macworkshops.com.          86400 IN      NS
mostsvr.macworkshops.com.
;; Query time: 7 msec
;; SERVER: 192.168.0.85#53(192.168.0.85)
;; WHEN: Thu Aug 18 08:13:52 2005
;; MSG SIZE rcvd: 72
```

DIG is shorthand for “Domain Internet Groper,” which, unlike its close cousin nslookup, pulls the actual raw DNS records from your server. Of utmost importance is the “NOERROR” message, which indicates the DNS goodness. Otherwise, you might see a “NXDOMAIN” message, or a “SEVERFAIL” message,

in which case you should go back, check your entries, then start and stop the DNS service. Next, we want to make sure that the reverse lookup’s working well, so we do:

```
; <<>> DiG 9.2.2 <<>> -x 192.168.0.85
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 41041
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1,
ADDITIONAL: 1
;; QUESTION SECTION:
85.0.168.192.in-addr.arpa.      IN      PTR
;; ANSWER SECTION:
85.0.168.192.in-addr.arpa.  86400 IN      PTR
mostsvr.macworkshops.com.
;; AUTHORITY SECTION:
0.168.192.in-addr.arpa.      86400 IN      NS
mostsvr.macworkshops.com.
;; ADDITIONAL SECTION:
mostsvr.macworkshops.com.  86400 IN      A
192.168.0.85
;; Query time: 2 msec
;; SERVER: 192.168.0.85#53(192.168.0.85)
;; WHEN: Thu Aug 18 08:19:53 2005
;; MSG SIZE rcvd: 111
```

to make sure that the reverse lookup record’s kosher as well. If you’re testing this on a machine other than the server, and don’t necessarily want to ssh into the server, you can simply add @ipaddress [hostname] to the dig statement to query a DNS server that you’re not currently configured to use, like so: dig mostsvr.macworkshops.com @192.168.0.85.

Now, you’re ready to promote your server to an Open Directory master. Go ahead and change the role in Server Admin, Open Directory. The following configuration window will pop up with a few fields to fill out:



**Get MacTech delivered to  
your door at a price **FAR**  
**BELOW** the newstand price.  
And, it’s **RISK FREE!****

**[store.mactech.com/riskfree](http://store.mactech.com/riskfree)**

From our brain  
to your brain...



**NEW!**



## SpiderWorks Books

Timely, quality content from respected authors at a great price. Available in softcover print editions and eBooks.



Download **FREE** Previews at [SpiderWorks.com](http://SpiderWorks.com)

**Create a new Open Directory master domain**

Creating a new Open Directory master domain requires you to create a new administrator account for that domain. This account needs to have a unique name, short name and user ID.

**New Account**

Name:

Short Name:  User ID:

Password:

Verify:

**Domain Info**

Kerberos Realm:

Search Base:

Search base is optional.

**Figure 4. Create New Open Directory master.**

Look at the bottom of Figure 4. Notice how important a correctly configured FQDN (Fully Qualified Domain Name) is to both the Kerberos Process as well as the search base for the LDAP directory. DNS and LDAP are inseparable; you cannot run LDAP “right” unless all of your DNS stars are properly aligned. Also new is the enforcement of a former best practice—creating a separate admin account for the new Directory. The local admin account (uid 501) you might have used to administer an Open Directory master under Panther Server, has absolutely no rights to the OD mater in Tiger Server. However, it’s easy to authenticate once, saving the new admin’s password in the keychain of the local account you might use to log into the Server, achieving virtually the same effect, though in reality, it’s an important conceptual difference, as we’ll see later on.

## Installing PHPLDAPadmin

Like the SquirrelMail ([squirrelmail.sourceforge.net](http://squirrelmail.sourceforge.net)) PHP-based web mail solution and the PHPMyAdmin ([phpmyadmin.sourceforge.net](http://phpmyadmin.sourceforge.net)) PHP-based MySQL Management tool, PHPLDAPadmin stands out as an analogue to these other two excellent applications. Being web-based, (especially using PHP) virtually guarantees that any OS with a somewhat current browser would be able to use it without compatibility headaches, as all of the data preprocessing happens on the web server. The first part of installing PHPLDAPadmin is a snap:

1. Download the script from [phpmyadmin.sourceforge.net](http://phpmyadmin.sourceforge.net)
2. If you have multiple sites set up on your server already, create a new virtual domain along the lines of “ldap.mydomain.com” and install the files there
3. Enable the php4\_module in the Web (service) > settings > modules section of Server Admin as in the figure below:

|                                     |                  |                 |       |
|-------------------------------------|------------------|-----------------|-------|
| <input type="checkbox"/>            | perl_module      | mod_perl.c      | libex |
| <input checked="" type="checkbox"/> | php4_module      | mod_php4.c      | libex |
| <input checked="" type="checkbox"/> | hfs_apple_module | mod_hfs_apple.c | libex |

**Figure 5. Activate php4\_module**

4. Next, it’s time to edit the config.php file for the PHPLDAPadmin. But first, it’s important to make a mental note—it’s OK to use an unencrypted web connection when the web-based form is running on the server it’s going to be connecting to, but if you want install PHPLDAPadmin on a server then connect to an LDAP directory on a different server, then SSL needs to be configured to make sure that the administrator credentials aren’t sent in clear text form over the “wire.” For simplicity, we’re going to use the configuration where PHPLDAPadmin is installed on the Open Directory Master itself.
5. First, navigate to the directory where you’ve unzipped or unpacked the PHPLDAPadmin script and locate the file called config.php.example, and make a copy, renaming it to config.php. This is the file you’ll use to configure the connection and authentication to your Open Directory Server.




**Figure 6. Copying Config File.**

Open up your config.php file in your favorite text editor. These days, it’s the free Textrangler 2.1 that’s floating my boat, along with it’s commercial counterpart, BBedit. First, we have to give our configuration a name, so change the following default setting:

```
$servers[$i]['name'] = 'My LDAP Server'; /* A convenient
name that will appear in the tree viewer and
throughout PHPLDAPadmin to identify this LDAP
server to users. */
```

Go ahead and enter a name between the single quotes, it’s not a DNS name, just a label.

Next, we need to change the “host” setting so that the script can connect to the server....

© 2004 Microsoft Corporation. All rights reserved. Microsoft, Your potential. Our passion., Entourage, and PowerPoint are registered trademarks of Microsoft Corporation in the United States and/or other countries. Apple, Mac, and Macintosh are registered trademarks of Apple Computer, Inc.

Microsoft®

Your potential. Our passion.™



## Get in touch with your inner suit.

Remain a free thinker while the other you gains a reputation as a serious go-getter. Microsoft® Office 2004 for Mac has intelligent new features that help you – or rather, the other you – keep tabs on every project, every e-mail, and every note. A smarter way to work. A smarter new Office. Go to [officeformac.com](http://officeformac.com) for a free trial.



P PowerPoint E Entourage W Word H Excel

Microsoft®  
**Office:mac**  
2004

```
$servers[$i]['host'] = 'mostsvr.macworkshops.com'; /*
Examples:

'ldap.example.com',
```

Now, configure the base DN (Distinguished Name), sometimes referred to as the “Search Base.” Because we’re going to be working with full DNs, we’ll make this a blank value.

```
$servers[$i]['base'] = ''; /* The base DN of your LDAP
server. Leave this blank to have PHPLDAPadmin
auto-detect it for you. */
```

Next we have to tell PHPLDAPadmin now we’re going to handle authentication. The most expedient way is to use the “session” method which relies on Apache:

```
$servers[$i]['auth_type'] = 'session'; /* Three
options for auth_type:
```

Like we did with the clearing out the base DN value, let’s do the same with ‘login\_dn’ and ‘login\_pass.’ Go ahead and save your edits.

Now, we’re ready to look at our Open Directory Master from the inside-out. Go to a web browser and type <http://yourwebsiteURL/ldap>. If your website was configured correctly, you should see the home page of PHPLDAPadmin:



Figure 7. PHPLDAPadmin Welcome Screen

## Getting in the Back Door

OK, so now it’s time to ask the Tiger Server Open Directory master to open up and say “aaaaah,” and let us inside. However, you can just go to the door and say knock knock, and when the voice inside says “who’s there?” answer with “admin uid=501.” In Tiger Server the local admin that installed the OS has no rights to the Open Directory master. So you’ll need to use the name “ldapmin uid=1000.” But even that’s not enough. You have to announce yourself using your full distinguished name, not just your short name and password. A DN or “Distinguished Name,” is basically your “long” LDAP identity, with contains the full path to where that identity lives in the LDAP Directory, along with, of course, the password necessary to authenticate so that you may do the business admins do.

Your DN would be something like:  
uid=ldapmin,cn=users,dc=nagitest,dc=macworkshops,  
dc=com  
And your password, would be, of course, your  
password: \*\*\*\*\*

If everything was configured copasetically, then you’ll see the promised land, which consists of this:



Figure 8. Successful Login to the LDAP Server

It doesn’t take a whole lot of figuring out at this point that it’s the plus signs that expand or collapse the LDAP view hierarchically, just as in Windows Explorer or Novell Console One. Now, some of the new features of Open Directory 3 become starkly apparent, we don’t even need to read the Server Documentation (although you really, really, should for a deeper understanding of Open Directory and *exactly* what its specifications are. After all, it’s *open*, as in Open-Source, and *open*, as in the sense that all of the directory data is presented in that documentation. Here’s what the first level of an OD master looks like at first glance:




Figure 9. Top Level of an OD Master.



DUCK.



DUCK.



DUCK.



GOOSE!

WORKING THROUGH  
SURPRISES REQUIRES  
THINKING DIFFERENTLY.  
WE CAN HELP.  
DON'T LET YOUR GOOSE  
GET COOKED...LOOK TO US  
FOR THE STRAIGHT STORY ON  
HOW TO MAKE THE MOST  
OF YOUR WINDOWS NETWORK,  
EXCHANGE & OUTLOOK  
EMAIL ISSUES, AND  
SECURITY CONCERNS—  
WITHOUT ANY  
MICROSOFT SPIN.

[WINDOWSITPRO.COM](http://WINDOWSITPRO.COM)

**WINDOWS IT PRO.  
THE LARGEST  
INDEPENDENT  
WINDOWS  
COMMUNITY  
IN THE WORLD.**

---

**Windows**ITPro  
CONNECTING THE IT COMMUNITY

What immediately leaps out is the first addition to Open Directory in Tiger Server, the “accesscontrols.” Other additions include “certificateauthorities and filemakerservers.” We should know by now that the ACLs for HFS on OS X Server are a *huge* leap forward.

Open Directory’s slapd process loads several schema files, which live at:

```
/etc/openldap/schema/core.schema
/etc/openldap/schema/cosine.schema
/etc/openldap/schema/nis.schema
/etc/openldap/schema/inetorgperson.schema
/etc/openldap/schema/misc.schema
/etc/openldap/schema/samba.schema
/etc/openldap/schema/fmserver.schema
/etc/openldap/schema/apple.schema
```

Integrating OS X into complex networks will never be the same again, as the answer to many questions is now “Yes, OS X Server can do that,” rather than “No, it doesn’t support that.” And just when we were getting used to understanding the limitations of those POSIX permissions, and working hard to transpose that understanding to the new ACLs for files and folders, now we find, squirreled away inside the “black box,” a new form of ACLs, sometimes referred to as DACs, or “directory access controls.” Let’s take a look at the default access controls, which live at the following distinguished name:

```
cn=default,cn=accesscontrols,dc=nagitest,dc=
macworkshops,dc=com
```

Herein lives an attribute called apple-acl-entry, and there are four default entries which are (please note that the backslashes indicate line breaks) numbered, much like firewall (IPFW) rules. Like many of the enhancements in Open Directory, the accesscontrols are a standard component of OpenLDAP:

```
1000:access to attr=userPassword by self write by
sockurl="ldapi://%2Fvar%2Frun%2Fldap" \
write by
group/posixGroup/memberUid="cn=admin,cn=groups,dc=nagitest,dc=
macworkshops,dc=com" \
write by * read

1100:access to attr=apple-user-authenticationhint by self
write by sockurl=\
"ldapi://%2Fvar%2Frun%2Fldap" write by
group/posixGroup/memberUid="cn=admin,cn=groups,\
dc=nagitest,dc=macworkshops,dc=com" write by * read

1200:access to attr=apple-user-picture by self write by
sockurl=\
"ldapi://%2Fvar%2Frun%2Fldap" write by
group/posixGroup/memberUid="cn=admin,cn=groups,\
dc=nagitest,dc=macworkshops,dc=com" write by * read

1999:access to * by sockurl="ldapi://%2Fvar%2Frun%2Fldap"
write by\
group/posixGroup/memberUid="cn=admin,cn=groups,dc=nagitest,dc=
macworkshops,dc=com" write by * read
```

While not a whole lot of practical lore (or documented experience) exists with regards to the capabilities of these DACs and Tiger Server, there’s always the OpenLDAP.org mailing lists and forums, although it’s not hard to imagine how they could be

applied, expanded, contracted or utilized to open up or close of certain areas of an LDAP directory. Apple’s Open Directory documentation (<http://www.apple.com/server/documentation>) characterizes the DACs this way: “Open Directory provides the ability to define directory access controls (DACs) to all parts of the LDAP directory, providing fine-grained control of who has permission to modify what. Open Directory stores the DACs in an apple-acl record that you can edit using the Inspector in Workgroup Manager.” Although the documentation advises admins to use Workgroup Manager to modify DACs, it is also easy to use PHPLDAPadmin as well, although the long-term consequences of using a third-party tool to manage an Open Directory master aren’t widely known. OS X Server admins are familiar with the dialog that appears when creating a new administrator account for a Directory Service domain, asking if the admin has right to modify users, groups and computers lists, along with associated managed preferences. Now with DACs, there’s a tool available to limit access for admins on an attribute by attribute basis, though no GUI exists for it in yet in Workgroup Manager.

Personally, I’d feel more comfortable using a tool with a hierarchical view, rather than the flat-list of attributes presented by Workgroup manager, but I’m a visual person. Many of us who’ve worked with Panther Server remember the dialog that appears when creating a new admin account asking if the new admin has rights to edit user accounts and preferences, as well as the same for groups and computer lists, but DACs allow for even finer tuning. For example, tinkering with DAC 1000 may affect the ability of admins to access the LDAP process itself. It also appears that the DACs are applied in ascending order from the lowest number to the highest number, suggesting that it might be a good idea to carefully consider how any DACs you’d add to the list might interact with others, or if it’s even a good idea to modify the default list at all, and if it is modified, how and where to modify.

## The World is Not Flat

One of the highly touted features of Open Directory 3 (again from OpenLDAP) is schema replication. An LDAP schema consists of a collection of attributes and organizational structures within a directory. Like Novell’s eDirectory and Microsoft’s Active Directory, Open Directory has the capability to scale to encompass the scope of a company, educational institution or governmental department that might be spread out over a large geographical area, with multiple branch offices. However, the default schema that shipped with Panther Server and now with Tiger Server assumes a flat directory with a single organizational unit, as if every single employee, computer, and department existed at corporate headquarters. It’s always been possible to create organizational units within Open Directory, even with Panther Server, except for one small problem. Even with a flat directory structure, large organizations need directory replicas to enable authentication at remote locations or as backups of the directory, and slurpd, the replication process of

Connect your people, contacts, relationships, projects, tasks, documents, blogs, web content ... your knowledge itself ... and magic follows.


Most businesses already have this data but in different systems, used by different people, in different departments. **What a nuisance!** NetTeam Server brings it all together, in a web app that everyone can use, anywhere. **What a difference!**

## NetTeam Server

NetTeam Server is a business process, content management, collaboration and social networking web app for businesses and organizations of any size. It offers people, project, task and document management services and has a powerful API to support customization.

The triangle represents NetTeam Server's unique combination of functionality for three critical areas: **Process, Content and Community**. We're a good choice even if you only need one of these, but if your business operations embrace two or three, you'll love what we can do for you.

- Project Workspace
- Team
- Tasks
- Documents



- Web Content/Site Management
- Document Library
- Project and Shared Blogs
- News Editor

- Users, roles and relationships
- Simple CRM
- Social Networking
- Profile & Personal Blogs

## Features

NetTeam Server's fundamental constructs are People, Projects, Tasks and Documents.

These four are central to all business activities, so we bring them together in a coherent workspace that makes NetTeam Server a true *Business Operating System*. The portal interface can be tailored to match client branding and linked systems, and includes five Editors (see screenshot) and a modern, AJAX-enhanced, configurable user interface.

User roles determine access privileges and which (if any) tools are presented on login. Blogs are used extensively to support publishing, information and knowledge management. Wikis will be available in a late-summer update.



NetTeam Server is available for Mac OS X Server, Linux and Windows platforms and supports all leading web browsers. A Web Services API allows tight integration with other systems and single sign-on. We also offer a Java mobile client which can be tailored to support mobile workforce applications.

It has been proven in deployments serving over 10,000 users and version 2 launches at the Collaborative Technology Conference in Boston, June 19-22 ([www.ctc2006.com](http://www.ctc2006.com)).

OpenLDAP, didn't support replicating customizations to the schema in Panther Server.

Replicas are attached to masters as read-only copies. The master slapd (OpenLDAP) process writes incoming changes to the directory to a file, which is monitored on a continual basis by slurpd, which reads the file, then updates the replicas with the ldap network protocol. The file slapd writes is an LDIF file (LDAP information file), which is also the LDAP import/export format. It then writes an LDIF version of the change to what is called a "replication log." So, even if you could neatly divvy up your directory so that certain users, groups, and computers lived in a different OU (organizational units) for different departments, your replication process wouldn't reflect that, or might even simply fail.

With Tiger Server, that's all changed, as least as far as schema replication support is concerned. Unfortunately, Apple's administration tools (Workgroup Manager, Server Admin) still operate under the assumption that the directory they read and write to is one large, flat space. Apple's directory service lookupd daemon is also painfully unaware of changes to the directory, and only finds the attributes that the mappings of the LDAPv3 Directory Access plugin allows, making it somewhat difficult to administer customized directories, or even create those customizations. PHPLDAPadmin, it turns out, is a great way to do so.

## Create An OU

Before testing schema replication, you'll need to have an Open Directory Master set up, and another installation of Tiger

Server to use as a replica. Keep in mind you'll also need to have two distinct licenses, as each Server uses the serial number as part of the ssh authentication necessary to start the replication process. Establishing a replica's a cinch, just enter the ip address of the OD master in the Open Directory settings in Server Admin, enter the root password on the OD master, the short name of the directory admin and its password, then wait until the process completes.

Let's say we were sysadmins for a school that wanted to create three OUs: faculty, staff, and students. We can use PHPLDAPadmin to do so; templates for common LDAP objects like OUs are included! After logging into PHPLDAPadmin, we can then go ahead and create our OUs by clicking on the "gold star" at the bottom of the directory tree:




Figure 10. Create New LDAP Entry

What an array of spiffy LDAP objects to choose from! Some, obviously, are going to be more useful to us than others, but I already am thinking of uses for the Samba 3 User and Samba 3 Group Mapping objects for mixed networks. Also intriguing is the "custom" template where it would be possible to store arbitrary attributes and retrieve them at will, almost like metadata. But we're after OUs, so create three of them: faculty, staff, and students:

# Problems with fingerprints, dust and grime?

- Safeguard your LCD display from dust and scratches
- Enhances visual image
- Set onto LCD itself
- An affordable yet beneficial investment that may extend the life of your screen

# Screen Protector For MacBook Pro

Singapore  
JCHS Media Pte Ltd 47 Beach Road, #04-04, Kheng Chiu Building, Singapore 189683 Tel: 65-6334-6618 Fax: 65-6334-6645 Corporate Website: [www.jchsmedia.com](http://www.jchsmedia.com)

www.mymobilejuice.com

Designed by: Justin Seah



## Lasso Professional Server 8.5

### Universal Binary for Intel Macs:

Native support for the best speed and compatibility.

### Includes AJAX Tags and Functions:

Supports Web 2.0 techniques for modern, dynamic Web design.

### Built-in LDAP Integration:

Integrate Lasso with office data and login systems.

### Connect to Any Database:

Including Oracle, PostgreSQL, MySQL, FileMaker, ODBC, and JDBC.

### Improved DNS, Email, iCal Support:

Extending Lasso's support for Internet protocols and open standards.

### 4 FREE well commented Lasso solutions to get you started:

**Blog:** Teaches beginners and advanced users alike how to create an online diary, or web log.

**QuickPoll:** Add a poll to your Web site with a question and answers, cast a vote or see results.

**Collaborate:** Enables knowledge sharing among a group of developers and administrators to streamline Web site development from remote locations or geographically dispersed teams.

**Message Board:** Source code for the Message Board solution is completely open and expandable to suit your needs.

### SPECIAL OFFER FOR MacTECH READERS

Try Lasso Professional Server 8.5  
FREE for 30 days.

[www.TryLasso.com](http://www.TryLasso.com)  
[lassosales@omnipilot.com](mailto:lassosales@omnipilot.com)  
800-678-9958

Then get 60 days of FREE  
support with your purchase.






Figure 11. Create LDAP Object

Next, log into Workgroup Manager on the OD master and create a new admin account with full rights called “facultymin.” You’ll notice that, like all users accounts, facultymin winds up in cn=users. Now, we can use PHPLDAPadmin to move facultymin into the faculty OU. If we also want to store groups, computers, and other things, we might want to consider creating copies of other LDAP “folders” within our OU.




Figure 12. Organizational Units

Now, we’ll move the facultymin user into the faculty OU. The part that gets hairy here is that after the move, Workgroup Manager cannot see the user account anymore! However, since facultymin is a directory administrator, we can still log into PHPLDAPadmin using the full DN of the user (uid=facultymin,ou=faculty,dc=nagitest,dc=macworkshops,dc=com) and we even have read/write access. But as far as lookupd and Workgroup Manager are concerned, the facultymin user’s gone.

Now it’s time to see the schema replication in action, and because an LDAP replica is read-only, simply use the anonymous bind setting with PHPLDAPadmin to log into your replica (you’ll need to install it there too) and check to see if the schema replication’s working. If everything’s copasetic, your OUs and facultymin user should now appear, almost instantly, in the replica LDAP directory.




Figure 13. After the Move.

## OU-ch

For a real-world deployment, we’d want to figure out a way that Workgroup Manager could still edit user, group and computer accounts that had been moved into OUs, or better yet, a way to create accounts in the proper OUs in the first place—without having to use PHPLDAPadmin or some other script to move them around. We’d also have to carefully consider how to adjust the DACs so that the admins like our facultymin user only could write to the objects in their own OU. Other important considerations would be how this would affect workgroup management, and cross-platform directory service integrations with the Active Directory or LDAPv3 plug-in, and access to directory information by other service like Apache and FTP. With so many open questions, I can only say: stay tuned for part deux.

## In Next Month’s Source Hound

Part deux of knock, knock, knocking on LDAP’s door. Hopefully, we won’t have to break it down, although we’re going to breakdown how many of the above Open Directory questions we can realistically answer, and see just practical it would be to administer an Open Directory deployment with Organizational Units (OUs) and Directory Access Controls (DACs).



## About The Author



*Dean Shavit is an ACSA (Apple Certified System Administrator) who loves to use a Mac, but hates paying for software. So each month he’s on the hunt for the best Open-Source and freeware solutions for OS X. Besides surfing for hours, following the scent of great source code, he’s a partner at MOST Training & Consulting in Chicago, where he trains system administrators in OS X and OS X Server, facilitates Mac upgrade projects for customers, and writes for his own website, [www.themachelpdesk.com](http://www.themachelpdesk.com). Recently, he became the surprised father of an application: Mac HelpMate, available at [www.machelpmate.com](http://www.machelpmate.com). If you have questions or comments you can contact him: [dean@macworkshops.com](mailto:dean@macworkshops.com).*

Visit [store.mactech.com/riskfree](http://store.mactech.com/riskfree) to order a One Year Subscription

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# AN OPEN SOURCE PRIMER

## A PRACTICAL GUIDE TO USING OPEN SOURCE SOFTWARE (OSS) ON MAC OS X FOR THE NON-DEVELOPER

BY EMMANUEL STEIN

### Introduction

Readers of Dean Shavit's "The Source Hound" and Ed Marczak's "Mac In The Shell", are no doubt familiar with open source software (OSS). With Darwin as OS X's core, many if not most of the Mac OS's most critical components are based on, if not entirely built upon, open source projects. To exploit the capabilities of OS X it is, therefore, critical to have a good grasp of OSS and the many ways you can leverage open source solutions on the Mac.

For Mac aficionados without a UNIX background, the world of open source can be rather daunting at first. Given all the distribution formats out there and the occasional need to compile software, it is not surprising that many Mac Users have not exploited OSS beyond what comes with OS X. Even non-developers who know how to compile their software, often run into issues with source code that requires special compiler options to successfully compile, or for which they may need to edit the Makefile and the like. Dependencies are also a source of problems and are often at the root of problems getting OSS to run properly. Finally, there is the issue of source code not ported to the Mac OS or Darwin, for which one requires a cross compiler. Source code not ported

to the Mac OS regularly stumps many users who do not realize that source code is frequently machine dependant in nature and therefore may not work on all platforms.

There is hope, however. Package management systems like Fink and Open Darwin Ports (think Apple Software Update for OSS) offer automatic downloading, compilation and linking, as well as updating. Although these tools are wonderful and make acquiring and deploying OSS a breeze (even for GUI diehards), they do not always include every piece of OSS you may want or need. Also, newer projects may not even have packages available, much less be indexed by the extant package management systems. In such situations, many people may give up and possibly look for a "safe" commercial alternative. I hope to change that with this primer by going over the convenient package management tools available, and going over the simple steps for compilation. Further, I will show you how to identify and where to find OSS that has been ported to the Mac.

Though, it is beyond the scope of a primer to go into porting OSS to the Mac, I will nevertheless touch on the conceptual issues involved. If nothing else, it will enable readers to better understand why certain code does not compile on the Mac and for more adventurous readers, shows where to find the porting instructions in a source distribution.

If you are a developer or UNIX geek, you will likely be familiar with the subjects covered. However, for the rest of us, get ready to enter the wonderful world of OSS. Far from being a "developer-only" community, the rich and vibrant world of OSS is chock full of cutting edge and highly useful non-developer tools like GIMP (Graphic Image Manipulation Program) and Blender (an advanced 3D modeling tool). Moreover, with the advent of OS X and object oriented desktop environments like Gnome and KDE,

# KNOW WHO'S WATCHING YOU.

**PROTECT YOUR PRIVACY  
MONITOR COMMUNICATIONS  
SECURELY ERASE FILES  
BLOCK ANNOYING WEB ADS**

**SUPPORT  
FOR APPLE  
MAIL**



A Division of Smith Micro Software

**Available from your favorite retailers and catalogs. Download a FREE Trial at Allume.com!**

One hundred percent 10.4 "Tiger" compatible! Download a FREE trial version of Internet Cleanup at <http://www.allume.com/mac/cleanup/ today!>

©, TM, and © 2005 Allume Systems, Inc. Mac OS X 10.3 "Panther" is required for Internet Cleanup 3.0.

the OSS universe has become increasingly accessible to GUI folks. With every passing day, open source developers are innovating and facilitating methods of OSS distribution. Whether via an Aqua wrapper to X11-based applications, like Aqua Ethereum, or via graphical front-ends to package management tools, such as Fink's Fink Commander and Open Darwin Port's Port Authority, OSS developers are clearly making every effort to reach the Mac audience.

## What Exactly is Open Source Software Anyway?

Open source, is typically used to refer to non-proprietary software, distributed with source code and a "copyleft" style license, allowing anyone to add new features to, or improve the source code as they see fit. Open source is however, more than a way of distributing software. It is a way of life with a deeply evolved philosophy that is all about fun and exploration in an intellectual arena. Central to that philosophy is the notion of hacking and an evolved spirit of play that goes beyond concepts of work and survival, to paraphrase Eric S. Raymond, a prominent OSS advocate and one of the few to have successfully hacked the three major open source UNIX projects: LINUX, BSD, and GNU. There are also practical characteristics specific to OSS that go beyond the qualification of providing the source code in a software distribution. Specifically, to "officially" qualify as open source, software must meet a set of well-defined guidelines set forth in the several extant open source licensing schemes. These include the BSD artistic License, the X11 license, GPL and other "compatible" license formats (see <http://www.gnu.org/licenses/license-list.html>). Though these licensing schemes differ slightly they are united by a similar philosophy that is delineated by the Open Software Initiative's (OSI) Open Source Definition (OSD) v1.9, whose axioms are available at (<http://www.opensource.org/docs/definition.php>).

Conceptually, open source, also represents a new paradigm in software development. To use the terminology coined by Raymond, in his seminal treatise "The Cathedral and the Bazaar," OSS employs the Bazaar model of development while traditional, closed source software, represents the "Cathedral Model." Without going into the details of each development model and its respective benefits and weakness, I have outlined the characteristics of open source software below:

- Releases are made as frequently as possible. There is usually a stable and an unstable release, the later with more features, but in need of testing.
- Users of the software, whether developers or not, contribute to the development, documentation and distribution of the software. This participation is a key element in the growth and sustainability of the OSS movement and is a way of giving back. User participation varies from submitting bug reports to taking an active role in the development of a project.
- The source is made available with the intent of enabling users to hack the source for educational and practical purposes.
- Whenever possible it is best to avoid forking development (e.g. The split with Emacs and xemacs is a good example of this) and duplicating effort. The ultimate goal is to work together to extend existing projects and only introducing novel projects to fill a niche not already saturated.
- Design the program, when possible, to be easily ported to as many other platforms as possible and embrace modularization of code for facilitated distribution of development work.

## Prerequisites

The vanilla install of OS X is packed with so much OSS the list is too long to reproduce here. However, so as not to rattle users who may be unfamiliar with these tools, Apple, being Apple cleverly hid them from view, much as they did the Terminal application. Even every-day applications like Safari are based on open source projects. In fact, without OSS there would be no Mac OS X!

The point is, with a standard install you can use all the OSS bundled with the Mac OS and many outside projects that are specifically packaged for the Mac. These include the OSS listed on Apple's OS X downloads page in the "UNIX & Open Source" section. However, if you want to use any OSS that requires, god forbid, compilation, or software which relies on the X11 windowing system, you have to go beyond the vanilla install.

The following is a list of the pre-requisites required for many if not most of the OSS mentioned in this primer. Installing this recommended software will enable you to get the most out of what is available in the open source community and ensure that you do not run into dependency issues.

1. If you haven't done so already, install Apple's X11 implementation. This is a custom option for the standard Mac install and can be found among the packages located on the install DVD for Tiger under the

**“Someone finally did it...  
Website tracking the Mac way!”**

- Neil Ticktin, MacTech Magazine Editor-in-Chief  
*and VisiStat.com Customer*



System>Library>Packages directory, for those of you who have already installed OS X and just need to add that package.

- The Developer Tools CD contains a rich array of Apple modified OSS needed for compiling software, as well as, many of the dependencies upon which OSS projects rely. Although the standard install of Developer tools will be adequate for most needs, I recommend also installing the X11 SDK and the optional compiler packages.

Note: Once you have installed the developer tools, be sure to repair permissions, as they will be changed in the process of installation and, if not remedied, may adversely affect your system performance.

## Open Source the Easy Way With Fink and Darwin Ports

Both Fink and Darwin Ports are amazing package management systems designed specifically to enable Mac users to benefit from the diverse range of OSS without having to manually compile or port source code. In addition to having easy to use command line interfaces, both Fink and Darwin Ports have mature GUIs that match virtually every available terminal option (Figures 1-4).




Figure 1. Darwin Ports' PortBase Graphical installer and Updater




Figure 2. Darwin Ports' PortView




Figure 3. FinkCommander GUI

These environments are ideal for both novice and experienced users, who simply want the equivalent of Apple's Software Update for their OSS. Frankly, I use both all the time on my work machine to get my OSS fix. Although I have been known to compile the occasional program, I have gotten so used to these wonderful utilities that I hardly ever have to compile these days. Each package manager offers access to a bewildering array of software, with Darwin Ports having 2,292 packages and Fink having 5,013 packages across 23 categories! What's more, via the GUI, you can choose to install the binary or source versions for maximum flexibility. Even if you only use Fink or Darwin Ports for one piece of OSS software, it is worth the download and install. Both detect dependencies for your desired package and automatically download and install the support files for you. This dependency checking feature alone has saved me countless hours hunting down library files and widget kits needed to run a simple OSS application. Finally the distribution options for these package managers is phenomenal, with .dmg all in-one installers familiar to any OS X user and the ability to obtain the source directly via cvs. Choose one or



Our headsets brought the world closer to the moon.  
Now they can keep you close to your music.



**PULSAR™ 590A**  
ULTIMATE STEREO BLUETOOTH® HEADSET

**MUSIC AND MOBILE** - Listen to music or movies, and answer phone calls.

**TALK AND LISTEN LONGER** - Up to 12 hours talk time and 10 hours stereo listening.

**UNIVERSAL ADAPTABILITY** - Enjoy Bluetooth functionality with most devices.

In 1969, a Plantronics headset carried those first words from the moon, "That's one small step for man..." Our latest Bluetooth headset, the Plantronics Pulsar 590A, was designed for versatility here on Earth. Seamlessly switch between your Bluetooth phone and your favorite music, so you'll never miss a call. And with the Universal Adapter, you can experience Bluetooth stereo listening on most laptops, Macs, TVs, DVD players and MP3 players. So your next mission will be as enjoyable as ever.

We've been to space, now it's your turn. Enter to Win a Trip to Space—visit [Plantronics.com](http://Plantronics.com) or text "headset" to SPACE (77223).



**PLANTRONICS.**  
SOUND INNOVATION™

NO PURCHASE NECESSARY TO ENTER OR WIN IN THE PLANTRONICS TO SPACE AND BEYOND PROMOTION. Must be 18 or older and a legal resident of the U.S. or D.C. Game begins 12/23/06 at 9:00:01 AM CT and ends 12/31/06 at 5:00:01 PM CT. Void where prohibited. Subject to the Official Rules, available at [www.plantronics.com](http://www.plantronics.com). See Official Rules for how to enter and complete prize details. ©2006 Plantronics, Inc. All rights reserved. Plantronics, the logo design, Plantronics Sound Innovation and Pulsar are trademarks or registered trademarks of Plantronics, Inc. The Bluetooth name and the Bluetooth trademarks are owned by Bluetooth SIG, Inc. and are used by Plantronics, Inc. under license. All other trademarks are the property of their respective owners. Space photography courtesy of NASA.



both, but Fink and Darwin Ports are must-have additions to any Mac user's OSS toolkit.

- You can obtain Fink and the associated GUI, Fink Commander, at <http://Fink.sourceforge.net/download/index.php?phpLang=en>
- Darwin Ports may be found at <http://darwinports.opendarwin.org/> and the corresponding GUI tools, including Port Authority at <http://www.wordtech-software.com/dpgui.html>

## To Compile or Not To Compile...

Often OSS will come in a variety of package and binary formats. While many, given the choice, will instinctively go for the binary or packaged formats, there are certain benefits to compiling. These include, the ability to better monitor the installation, to effectuate custom configurations, or for reasons of security. Regarding the last point, it is rare that package maintainers are ever malicious and, provided you get your package or binary from a reputable source (e.g. sourceforge.org, freshmeat.net, are two popular examples), you risk little or nothing. Ditto, for users of package utilities like Fink and Open Darwin Ports. Users of these utilities can also compile from source.

Regardless of whether your preference is for pre-compiled and packaged binaries or source auto-compiled using Fink, there will come a day, mark my words, when you will need to compile a piece of OSS. It may be because the binaries were not updated for the latest OS update, or because there are no OS X packages out there and/or your package utility has not indexed the particular piece of OSS you need. While I know many people who would prefer to eat a bug than open the terminal and compile software, developers often make it very easy to compile their source code. Much of the time the following terminal command is all that is needed:

```
$ ./configure && make && make install
```

There are, nonetheless, instances in which this command will not suffice. We will discuss this further in the next few sections and offer reasons, as well as, tips for doing more advanced compiling. We will also cover what porting source to the Mac involves.

## Where To Find Mac Friendly Source Code

Although, sites specifically geared towards OSS on the Mac exist, they are often limited to pre-made packages and can be lacking in terms of breadth and quantity of OSS software. In contrast, platform independent sites like sourceforge.org and freshmeat.net not only offer Mac-specific source and binaries, but also offer a rich and centralized repository for the latest and greatest in OSS. I have seen many advertisements for CD distributions of

Mac compatible OSS binaries and source, but have always felt that these missed the whole point of OSS: To distribute up-to-date builds with frequent patching, that frankly only a medium like the Internet will allow.

## What's in the Source?

With all this talk of source code, you might be wondering what it is and how it is distributed. Even though compilation is usually accomplished with a simple three-step command (e.g. `configure && make && make install`), knowing what to look for in a source distribution can go a long way to ensuring a successful build of your OSS.

Your typical source distribution comes packaged as a compressed tarball (e.g. `mysourcecode.tar.gz`). When expanded and untarred, the source distribution will likely contain several files including `.c` and `.h` files, which represent the source code and header files, with `README`, `INSTALL`, and sometimes `PORT` files. Although most source code is written in a version of the C programming language, hence the `.c` files, other distributions may be written in perl, python, and a host of other languages. As such, you may encounter distributions without `.c` and `.h` files. However, Perl and Python do not require compilation, as they are interpreter-based scripting languages and are thereby much easier to deploy.

For easy compilation, developers often supply a configure script that is generated with the `autoconf` OSS utility. Alternatively, you may find that your OSS distribution uses the `xmkmf` script to invoke the `imake` program, which, in turn, will construct the make files needed for compilation. When confronted with OSS based on `imake`, be sure to read the `INSTALL` file for details. However, in most cases the following command can be employed to generate necessary files for compilation:

```
$ xmkmf -a
```

The compilation may then be completed by issuing the following terminal command:

```
make && make test && make install
```

There will also be cases in which no configure script is supplied with the source. You may, nevertheless, generate one by executing the `autoconf.sh` script. Please note that this applies only to source, which is based upon `autoconf`. Most commonly, you will run into this scenario when obtaining the source directly from a CVS (Control Versioning System) repository on the Net. As previously mentioned, you should then be able to run the `configure` command, followed by the `make` and `make install` commands. Though this is the typical manner in which compilation is effectuated, be sure to read the `INSTALL` file, as it should contain more precise install instructions. The `PORT` file, mentioned earlier offers suggestions for developers wishing to port the OSS to another UNIX platform, like OS X.

# No Limits

[www.MacSales.com](http://www.MacSales.com)



**Premier External  
Storage Solution**  
USB/FireWire Port  
Expansion Hub

**&**



Home Entertainment Center



Total Portability

**Up to  
750GB!**



On the Job



Complements All Computers



Maximize your mini



**Extra  
Ports!**

## miniStack™ V2

**newer technology**

Own the Future - **TODAY!**  
[macsales.com/storage](http://macsales.com/storage)

The NewerTech **miniStack™ V2** now with both rear and side FireWire and USB Hub ports. Now up to 750GB! Home Entertainment, Audio/Video, Backup, Bootable Data Storage, Music, Graphics, Performance RAID, & More! Storage Capacities from 80GB to 750GB.



**Other World Computing**

Serving the Computer Universe since 1988  
Visit [macsales.com/storage](http://macsales.com/storage) 800.275.4576

Check out our full line of external storage solutions plus much more at [macsales.com](http://macsales.com)



\*Test performed with 750GB model. Similar results with 80GB to 750GB models.

Cinema Display and Mac mini are trademarks of Apple Computer, Inc.

© 2004 Other World Computing, Newer Technology, a part of New Concepts Development Corp., 1004 Courtlands Dr., Woodstock, IL 60098

Prices, specifications, and availability are subject to change without notice. Items returned within 30 days may be subject to a restocking fee. No return will be accepted without Return Merchandise Authorization number.



| Name           | Date Modified         | Size  | Kind                   |
|----------------|-----------------------|-------|------------------------|
| acconfig.h     | Aug 26, 2004, 3:56 AM | 4 KB  | C Header Source File   |
| aclocal.m4     | Jan 13, 2005, 3:22 AM | 36 KB | Document               |
| AUTHORS        | Jan 11, 2005, 3:41 AM | 4 KB  | Document               |
| autom4te.cache | Jan 9, 2005, 8:01 AM  | ---   | Folder                 |
| ChangeLog      | Aug 26, 2004, 3:56 AM | 4 KB  | Document               |
| config.h.in    | Jan 11, 2005, 3:27 AM | 4 KB  | Document               |
| configure      | Jan 11, 2005, 3:46 AM | 1 KB  | Unix Executable File   |
| configure.in   | Jan 11, 2005, 3:46 AM | 8 KB  | Document               |
| configure.ac   | Aug 26, 2004, 3:56 AM | 1 KB  | Document               |
| COPYING        | Aug 26, 2004, 3:56 AM | 20 KB | Document               |
| curses.c       | Jan 11, 2005, 3:32 AM | 16 KB | C Source File          |
| depcomp        | Aug 26, 2004, 3:56 AM | 12 KB | Unix Executable File   |
| display.c      | Jan 11, 2005, 3:35 AM | 8 KB  | C Source File          |
| display.h      | Jan 11, 2005, 3:30 AM | 4 KB  | C Header Source File   |
| dns.c          | Jan 11, 2005, 3:32 AM | 36 KB | C Source File          |
| dns.h          | Jan 11, 2005, 3:30 AM | 4 KB  | C Header Source File   |
| FORMATS        | Aug 26, 2004, 3:56 AM | 4 KB  | Document               |
| getopt.c       | Jan 11, 2005, 3:33 AM | 24 KB | C Source File          |
| getopt.h       | Jan 12, 2005, 5:25 AM | 8 KB  | C Header Source File   |
| getopt1.c      | Aug 26, 2004, 3:56 AM | 8 KB  | C Source File          |
| glib.c         | Jan 11, 2005, 3:36 AM | 16 KB | C Source File          |
| img            | Jan 11, 2005, 3:39 AM | ---   | Folder                 |
| INSTALL        | Aug 26, 2004, 3:56 AM | 8 KB  | Document               |
| install-sh     | Aug 26, 2004, 3:56 AM | 8 KB  | Unix Executable File   |
| logfile        | Aug 26, 2004, 3:56 AM | 20 KB | Document               |
| Makefile       | Jan 11, 2005, 3:39 AM | 4 KB  | Document               |
| Makefile.am    | Jan 11, 2005, 3:10 AM | 4 KB  | Document               |
| Makefile.dist  | Aug 26, 2004, 3:56 AM | 4 KB  | Installer distribution |
| Makefile.in    | Jan 11, 2005, 3:48 AM | 20 KB | Document               |
| missing        | Aug 26, 2004, 3:56 AM | 12 KB | Unix Executable File   |

**Figure 4. A view of a typical autoconf-based source distribution**

## What if the Source Code for The Software I want Isn't Mac Compatible?

Unless specifically designed to be platform independent, the source code has to have been ported to the Mac architecture in order to compile and run on OS X. Unfortunately, the Mac is not binary compatible with Linux, so even Linux packages ported to the PowerPC architecture are not usable. Mac OS X differs notably from Linux, SVR 4-based systems, and other Unix variants in its lack of support of the Executable and Linking Format (ELF). The binary format specific to OS X is Mach-O. This is why even PowerPC architecture specific source and binaries are a no go in OS X.

If the OSS you need is not available for OSX, consider using an emulator or configure a dual boot option with Yellow Dog Linux, for instance. Unless you are a developer and are willing to give up a significant amount of your free time to port the OSS, the options I mentioned are your only choice. On the plus side, however, is that with each passing day more and more OSS is being ported to OS X, which though not totally Open Source itself, has become one of the leading OSS platforms.

## A Call To Arms

Although OSS is free, characterizing it as simply "free software" misses the essence of the OSS philosophy and the hacker ethic from which it originated. The open source movement, along with the

Internet has enabled developers to adopt a radically different development model. This new model, termed the "Bazaar" by Raymond, has and continues to prove itself as a preferred model for software development and distribution and will continue to make inroads into virtually all areas of the technoverse. Built on a foundation of cooperation, group effort, and imbued with a spirit of play and intellectual curiosity, OSS represents a novel paradigm for the exchange of ideas and has the potential to fundamentally alter how we think of and use technology. Moreover, the detailed philosophies coming out of the open source movement offer a dramatic and compelling alternative to the traditional Protestant ethic, which values work for work's sake. The vision presented by the OSS movement is of a more evolved and egalitarian society, in which the joy of hacking transcends the Protestant work ethic. As part of the Mac community, it is our collective responsibility, to not only take from the rich array of OSS, but to give back as well. For developers the meaning here is very clear. However, non-developers are far from excluded and play a crucial role as software testers, technical writers, and distributors of OSS. I urge you all to take the plunge into the world of OSS. Together, we can take ownership of the technologies upon which we depend and really make a Jobsian "dent in the universe!"



## About The Author

*Emmanuel Stein has been an avid Mac user since 1984 and has honed his cross-platform skills while working at France Telecom, Time Magazine and Reed-Elsevier. He has recently started his own Mac-centric consulting company, MacVerse, which offers implementation, system administration and development services geared towards the enterprise market. As a diehard GNU/Linux geek, he enjoys hacking open source software and experimenting with new open source projects on OS X. You may reach him at [macverse@mac.com](mailto:macverse@mac.com)*

**MACTECH**  
M a g a z i n e

**Get MacTech delivered to  
your door at a price FAR  
BELOW the newstand price.  
And, it's RISK FREE!**

**[store.mactech.com/riskfree](http://store.mactech.com/riskfree)**

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# Deconstructing RSS 2.0

## Understanding How RSS Feeds Work

By Dave Woolridge

If you're one of the millions of people who maintain your own weblog, then odds are, you're probably already familiar with RSS or Atom feeds since most blog tools include support for offering your blog as a syndicated feed. If you don't write a blog, then you've undoubtedly seen the RSS icon displayed online (see **Figure 1**) or have been invited by web sites to subscribe to their free RSS feeds. In fact, it's such a hot technology these days that you would have had to have been living "off the Grid" for the last few years to have not heard about RSS.

With aggressive spam filters making e-mail communication difficult for even legitimate marketers and businesses, web feeds have become a safe and dependable method for you to successfully deliver news to your audience, as well as allow third-party sites to syndicate your feed content for expanded reach to new viewers. It's a win-win situation for everyone involved. Users can subscribe to only the feeds they wish to receive and third-party sites are provided with free content for their sites that ultimately drives additional traffic to your site via your feed's links.



**Figure 1. Typical RSS buttons seen on web sites and blogs, commonly referred to as "chicklets"**

While the RSS and Atom formats are both used throughout the Internet for content syndication, this article will focus on the RSS 2.0 specification. With support for multimedia enclosures and other multi-purpose features, RSS 2.0 is quickly becoming one of the most popular flavors for content syndication, providing site owners with a powerful vehicle for delivering a lot more than just news. Have you ever subscribed to a podcast? Yup, you guessed it... podcasts are RSS 2.0 feeds.

### How RSS Works

Before we dive in, let's dispel the most common misconception about web feeds. While many people use

the word "broadcast" to describe online content syndication, web feeds are not transmitted like radio or television signals. They are not beamed or sent to subscribers. A web feed is nothing more than an XML document that resides on a web server. This means that your news reader software or web browser is fetching the RSS feed from a specified URL, just like it would do to access an HTML web page. The reader/browser software then parses (translates) that XML document, providing it to you in a display format that's easy to read. When using a news reader application (such as NetNewsWire or NewsFire) or an RSS-savvy browser (such as Safari or Firefox), subscribing to an RSS feed is like adding a bookmark to your favorites list. The only difference is that an RSS feed will get automatically checked for new updates on a regular basis, which requires the reader/browser to go fetch the XML document from the feed's URL at each timed interval.

### XML Syntax

Since we'll be taking a look at the structure of an RSS feed and the specific functionality of each XML tag in the RSS specification, it's important to understand the basic syntax of XML (Extensible Markup Language). Like HTML code, XML consists of tags such as:

```
<title>My RSS Feed</title>
```

Unlike HTML, which helps define styles and formatting for text, XML tags strictly define the meaning and context of information, keeping all of the text neatly organized with tag names. You don't have to be a master of XML in order to write or modify your own RSS feeds, but there are a few basic rules that you should keep in mind.

- Every valid RSS feed needs to include `<?xml version="1.0" encoding="UTF-8"?>` as the very first tag at the top of the document. UTF-8 is what most feeds use as the text encoding and most RSS

## He's busy doing a software update on 1,037 computers.

Hard to tell though, isn't it? That's because he's automating it using FileWave® software. Totally cross-platform, totally automated. So, are you an IT administrator? Tired of manually installing OS updates, security updates and patches? Or maybe you just aren't keeping your software up to date? Automate it with FileWave® and save time for the important things.

**Hundreds of satisfied customers agree:  
FileWave® just works.**



*"We can do all updates via FileWave® with 99.9% success (you wouldn't believe me if I had said 100%)"*

— David Zubradt, FileWave Administrator for a leading US cable television corporation.

*"It's an indispensable product that allows me to support a district of our size with only two full-time staff members, yet it appears that there is an army of people working in the IS department."*

— Brent Haywood Urbandale Community School District, Urbandale, IA

parsers assume UTF-8 as the default if no encoding is specified, but if you need to use a different encoding for a special purpose, then change the encoding attribute accordingly.

- The nesting order of XML tags is very important in order for your code to be valid. For example, the following line is properly nested:

```
<skipDays><day>Sunday</day></skipDays>
```

while the same line of code below will cause errors since it contains invalid nesting:

```
<skipDays><day>Sunday</skipDays></day>
```

- XML is case-sensitive, so <skipdays> is not the same as <skipDays>. While some RSS readers and browsers may be smart enough to overlook case typos, you certainly want to avoid any potential problems with your feed, so if you're modifying your RSS code by hand in a text editor, make sure your tags conform to the RSS 2.0 specification.
- Be very careful when including HTML code within XML. If you need to include HTML code within the title or description tags of a news item, you need to either enclose all of the text within a CDATA block such as:

```
<description><![CDATA[My <b>Big</b> News]]></description>
```

or convert all HTML brackets and special characters into XML-safe entities such as:

```
<description>My &lt;b&gt;Big&lt;/b&gt; News</description>
```

## Reviewing an Example

To get a feel for what we are talking about here, let's take a look at an example RSS feed. RSS is short for "Really Simple Syndication," and as you'll see from the XML code in **Listing 1**, the format really is quite simple overall. RSS 2.0 feeds can be saved with a file extension of either .rss or .xml. **Listing 1** shows an RSS feed that includes two items. The first one is a typical news item, while the second one contains a media enclosure, similar to what you would find in a podcast.

### Listing 1: An RSS 2.0 Feed Example

```
<?xml version="1.0" encoding="UTF-8"?>
<rss version="2.0">
<channel>
  <title>SpiderWorks News</title>
  <link>http://www.spiderworks.com/</link>
  <description>Quality eBooks and Printed Books
from Respected Authors at a Great
Price!</description>
  <language>en</language>
  <copyright>Copyright 2005 SpiderWorks, LLC. All
```

## Big Security

for the small business



MacWorld "Best of Show" Winner,  
CRYPTOCARD introduces 5-User  
Kit for OS X, the "All-In-One-Box"  
Authentication Solution that lets  
your security grow with you.

**CRYPTO-Server 6.3** is the Mac-centric authentication solution for IT infrastructures organized around Mac OS X. Once implemented, your Mac becomes the centralized authentication and token management system for all users, regardless of their computing platform or location.

#### For Mac Desktops

- Enforced CRYPTOCARD authentication for desktop access – Startup, Screen Lock, Sleep
- Authenticated Fast User Switching
- Internet Connect MSCHAPv2 (Tiger)
- Cisco VPN Client for Mac Integration

#### Infrastructure Integration

- Open LDAP integration
- Open Directory integration
- Databases (MySQL, Oracle)

#### Network Integration

- RADIUS-enabled systems – VPN, Firewall, NAS
- Web Servers & On-Demand Computing – Apache, IIS, Citrix

#### MACTECH SPECIAL OFFER

**5-User Kit** for only **\$249**  
• includes 5 tokens, CRYPTO-Server software & 30-day support.

Email [mac@cryptocard.com](mailto:mac@cryptocard.com)  
and quote **SPMac249-01**.

*Promotion can end without notice.*

**CRYPTOCARD**  
Secure Password  
TECHNOLOGY  
[www.cryptocard.com](http://www.cryptocard.com)

```

rights reserved.</copyright>
<managingEditor>news@spiderworks.com</managingEditor>
<webMaster>support@spiderworks.com</webMaster>
<rating> </rating>
<pubDate>Mon, 24 Oct 2005 05:53:07 GMT</pubDate>
<lastBuildDate>Mon, 17 Oct 2005 08:13:02 GMT</lastBuildDate>
<category>
domain="http://www.spiderworks.com">Books</category>
<docs>http://blogs.law.harvard.edu/tech/rss</docs>
<generator>SpiderWorks Staff</generator>

<skipDays><day>Sunday</day></skipDays>
<skipHours><hour>14</hour></skipHours>
<ttl>30</ttl>
<cloud domain="http://www.spiderworks.com"
port="80" path="/rpc"
registerProcedure="rssPleaseNotify" protocol="XML-
RPC" />

<image>
<title>SpiderWorks</title>

<url>http://www.spiderworks.com/swbadge.gif</url>
<link>http://www.spiderworks.com/</link>
<width>120</width>
<height>35</height>
<description>Quality eBooks and Printed Books
from Respected Authors at a Great
Price!</description>
</image>

<textInput>
<title>Search SpiderWorks</title>
<description>Search the SpiderWorks Web
Site</description>
<name>q</name>
<link>http://www.spiderworks.com/search.pl</link>
</textInput>

<item>
<title>SpiderWorks Releases Danny Goodman's New
Dashboard Book</title>

<link>http://www.spiderworks.com/books/dashboard.php
</link>
<description>World-renowned JavaScript and
Dynamic HTML expert, Danny Goodman, shows you how
to build rock-solid, professional Dashboard widgets
for Mac OS X Tiger in his new book, Mac OS X
Technology Guide to Dashboard. Includes exclusive
widget debugging tool, The Evaluator! Available as
an eBook and printed edition at
SpiderWorks.com</description>
<category>
domain="http://www.spiderworks.com/topics/tiger.php"
>Mac OS X Technology Guides</category>
<source>
url="http://feeds.feedburner.com/spiderworks">Spider
Works</source>
<author>Dave Wooldridge</author>
<pubDate>Tue, 05 Jul 2005 2:37:01 GMT</pubDate>

<comments>http://spiderworks.blogspot.com/2005/07/da
shboard_comments.html</comments>
<guid>
isPermaLink="true">http://spiderworks.blogspot.com/2
005/07/dashboard.html</guid>
</item>

<item>
<title>Spiderworks Interview with Ben
Waldie</title>

<link>http://www.spiderworks.com/books/automator.php
</link>
<description>SpiderWorks recently sat down with
author Ben Waldie to discuss his new book, Mac OS X
Technology Guide to Automator. Listen to the full
interview online.</description>
<enclosure

```

```

url="http://www.spiderworks.com/audio/waldie.mp3"
length="243108" type="audio/mpeg"/>
<category>
domain="http://www.spiderworks.com/topics/tiger.php"
>Mac OS X Technology Guides</category>
<source>
url="http://feeds.feedburner.com/spiderworks">Spider
Works</source>
<author>Dave Wooldridge</author>
<pubDate>Mon, 02 May 2005 0:00:01 GMT</pubDate>
<comments>http://spiderworks.blogspot.com/2005/05/au
tomator_comments.html</comments>
<guid>
isPermaLink="true">http://spiderworks.blogspot.com/2
005/05/automator.html</guid>
</item>

</channel>
</rss>

```

At first glance, **Listing 1** looks like a typical XML document, but what sets it apart as RSS is the way the XML data is structured. The root element is, of course, the `rss` tag. Nested within that tag is the `channel` tag. It's within the `channel` tag that the meat of the document is stored. **Figure 2** breaks down the main ingredients of this RSS example into groups, providing you with an easy way to visualize the XML code in **Listing 1**.



**Figure 2.** The RSS example broken down into basic groups of elements.

The groups within the `channel` tag consist of two groups of data: (1) channel elements that describe details about the RSS feed itself, and (2) items that hold your news stories, podcast audio tracks, etc. While **Listing 1** and **Figure 2** only include two items for example purposes, you can add as many items as you want to your own RSS feed.

## Defining Your RSS Feed

As shown in **Figure 2**, before your actual news items are listed, your XML needs to include some basic information about the RSS feed itself. These tags are called channel elements and include important information like the title of your feed, the content's copyright, the date the feed was last updated, etc. Only a few of these elements are required, but the more information you provide, the more efficiently the

receiving RSS readers (known as aggregators) can handle and process your data. For example, including optional tags like `lastBuildDate` and `ttl` can help relieve the server load from your feed being requested unnecessarily since those tags specify when the feed was last updated and how long the data should be cached (stored temporarily) before refreshing with a new HTTP request of the feed.

While all of the channel elements are defined here, please refer to the code examples in **Listing 1** for the proper XML syntax of these tags.

### title

REQUIRED. The name of your feed, which is usually the same name of your blog or web site that's related to your feed.

### link

REQUIRED. The URL of your blog or web site (not the URL of your feed).

### description

REQUIRED. A very brief phrase or sentence that describes your feed's overall content.

### language

Optional. The language that the feed is written in. For list of possible language codes, please refer to: [http://blogs.law.harvard.edu/tech/stories/storyReader\\$15](http://blogs.law.harvard.edu/tech/stories/storyReader$15)

### copyright

Optional. The copyright notice for the feed's content. Do not use the actual copyright symbol since that special character may not display properly in RSS readers.

### managingEditor

Optional. The e-mail address of the editor of the content. This is not necessarily the author of the content, since the content may have come from multiple sources, but this is the person who is managing the feed.

### webMaster

Optional. The e-mail address of the webmaster who oversees all technical issues related to the feed.

### rating

Optional. The PICS rating of the feed, which helps adults control what online content is accessible by children. This tag is rarely used, but for more information on PICS, please visit: <http://www.w3.org/PICS/>

### pubDate

Optional. The publication date of the feed, which states the earliest date that the content can be publicly displayed. Most aggregators ignore this tag and instead focus on the `lastBuildDate` tag. The date should be formatted to conform to RFC 822, which can be found at: <http://asg.web.cmu.edu/rfc/rfc822.html>

# iAutomate.



Edit, optimize and convert all your images and video for optimal display on Web, Office apps, 3G phones, iPod and portable media players with a simple drag and drop.



**Save Time and \$50 on Equilibrium® DeBabelizer® Pro for Mac® OS X/Windows® XP when you order now.\***

\*For a limited time only, get a \$50 mail-in rebate when you order from your favorite software outlet. Go to [www.equilibrium.com](http://www.equilibrium.com) for details and call us at **866.EQUILIB** (+1.415.332.4343) with any questions.



# How to Stop Racing the Clock.



☐ Work Longer? ☐ Work Harder?

☒ Energize Your Mac!

**We know your day keeps getting longer and longer.** With every release of software, your Mac is bogged down even more. With every click, there's a pause. You find yourself working longer, working harder. *A faster Mac means that you can work faster, not harder - be more productive!*

**Let the original Mac Performance Shop help.** Daystar has been creating Mac speed for over 16 years. Whether your bottleneck is storage, connectivity or just raw CPU speed, we deliver the performance you need, where you need it.



**CPU Upgrades for Raw Speed.** We upgrade any Power Macintosh, any iMac Flat Panel, any PowerBook G3 and some PowerBook G4s.

**Fast and Large Storage for Real-Time Video.** Our **TURBOSATA** solutions can make your drives perform like RAM. Projects open in a flash and edit in real-time.

**Extreme Wireless.** Wireless is great, unless you're getting slow transfers. Even Airport Extreme's are slow when the signal is weak. Daystar can boost your signals and energize your wireless network.

**But, if You Really need a G5?** Daystar is the only Mac Performance Manufacturer that is also an Apple Authorized Reseller. Not only can you trade-in your system for the latest and greatest... but the Daystar Pro's can upgrade it for maximum performance!

**Call 877-439-8646 and beat the clock.**



Authorized Reseller



**Daystar Technology - Your Macintosh Performance Shop**

5018 Bristol Industrial Way, #202, Buford, GA 30518 USA  
Toll Free: 877-439-8646 or 770-614-5400

Daystar-Tech.com

Daystar-Forum.com

Daystar-Store.com

## lastBuildDate

Optional. The date that the feed was last updated. This is often one of the first tags that aggregators check to see if any new content was been added or updated since the last time the feed was requested. Like `pubDate`, this tag's date should be formatted to conform to RFC 822.

## category

Optional. If your blog or web site organizes blog entries and articles into specific categories, then this tag may help aggregators to categorize items accordingly. Unfortunately, there is no standard cataloging system, so often this tag only proves useful for your own site needs. The `domain` attribute typically refers to a URL for that category online, but if your site does not include unique web pages for each category, then you may want to just link to your home page.

## docs

Optional. This tag should link to the official RSS specification online. As an RSS 2.0 feed, this tag should point to: <http://blogs.law.harvard.edu/tech/rss>

## generator

Optional. If you used feed generator software to create your feed, then the application would give itself credit in this tag. For example, if you used FeedForAll to generate your feed, then this tag may read:

```
<generator>FeedForAll</generator>
```

## skipDays

Optional. This tag informs aggregators that the feed should not be read on certain days. Within the `skipDays` tag is a nested `day` sub-element, so that you can include more than one day within `skipDays`. Acceptable day values are: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, or Sunday. For example, to list both Sunday and Monday as days to skip, you would use the following XML syntax:

```
<skipDays><day>Sunday</day><day>Monday</day></skipDays>
```

## skipHours

Optional. This tag informs aggregators that the feed should not be read during certain hours. The `skipHours` tag syntax works exactly like the `skipDays` tag, except that the sub-element is `hour` instead of `day`. An acceptable hour value is any whole number between 0 and 23. Like `skipDays` and its `day` sub-element, `skipHours` can include multiple hour sub-elements.

## ttl

Optional. This tag represents the “time to live” with the value being in minutes. It tells aggregators how long the feed content should be cached on their end before

refreshing with a new HTTP request of the feed. For example, if you use 60 as the value, then aggregators will know that they need to wait 60 minutes before requesting a fresh copy of the feed. This can help alleviate some of your server load since it will decrease the number of redundant feed requests.

## cloud

Optional. This is probably the most rarely used and most confusing tag in the RSS 2.0 specification. Many developers who encounter this tag either don't understand how it works or have not figured out how to best utilize it. This tag represents a lightweight publish and subscribe feature that includes five essential attributes: `domain`, `port`, `path`, `registerProcedure`, and `protocol`. It is a way for RSS to leverage the power of web services like SOAP or XML-RPC to conserve server bandwidth, allowing aggregators to register to be automatically notified of any updates made to the feed. If you're not using SOAP or XML-RPC web services on your server, then you won't have a need for this tag. For more information on using SOAP with the optional `cloud` tag, visit: <http://blogs.law.harvard.edu/tech/soapMeetsRss>

## image

Optional. You may have noticed that some aggregators display a logo or badge from the feed they are displaying. They get this data from the `image` element, which includes six sub-elements: `title`, `url`, `link`, `description`, `width`, and `height`. For example, if you publish a sports news RSS feed called “Primo Sports Plus” which has its own unique logo, you can supply aggregators with that logo to help brand your syndicated content. There's no guarantee that aggregators will use the logo, but the extra marketing potential makes it worth including. `title` represents the ALT tag if the image is rendered in HTML and `link` represents the URL of your site if someone clicks on the image. The `url` attribute should be the direct URL to the actual image file itself. The image's `description` value is usually the same as your feed's `description` tag. It's important to note that the maximum image width value is 144 (defaults to 88 if unspecified) and the maximum image height value is 400 (defaults to 31 if unspecified). See **Listing 1** for the XML code syntax of the `image` tag and its nested sub-elements.

## textInput

Optional. This rarely used tag is ignored by most aggregators. It provides a way for aggregators to submit a search (if your site has a search engine) or submit user comments (if your blog accepts user input). If you decide to include this tag in your feed, it requires four sub-elements: `title`, `description`, `name`, and

# Intego Personal Backup X4

## Back up and protect your digital life

### Toolbar

Choose the operation you want to perform: backup, restore, clone, synchronise, move, and archive your data.

### Source

Select what you want to back up by simply dragging and dropping it onto the Personal Backup X4 interface.

### Status information

Display the progress of your backup, as well as information about the files being transferred.

### Scrolling display

View the selected backup script option in the scrolling display window.

### Statistics and graphs

See the speed of your data transfer in real time.

### Script creation

Create and edit backup scripts easily and use multiple options and exceptions.


### Destination

Back up your data to hard disks, CDs, DVDs, iPods, USB drives, iDisks, etc.

### Quick launch

Start and stop your backup task with a single click.

**NEW  
VERSION**



**For Macs  
with PowerPC  
or Intel processors**



### Intego Widgets

Discover the 3 new Intego Widgets informing you of the status of your protection, available updates and scheduled events.




### Compatible with iPod

Easily transfer and back up your data to your iPod.



### Integration with iCal

Schedule regular or one-time backups of your data and view them in Apple's iCal.



### Your updates with NetUpdate X4

Easily schedule, check and install new updates for Personal Backup X4, as well as updates for all Intego software installed on your Mac.



### Main features of Personal Backup X4

- Full or incremental backups
- Fast, reliable restores
- Startup volume can be cloned
- Data synchronisation
- Fully redesigned and resizable interface
- Archive encryption (AES 128)
- New menu integrating all Intego X4 software
- Personal Backup X4 Widget and Intego Widget
- Predefined backup scripts for music, video and more
- Scripts can be linked to one another
- Exception list
- Data copied in its original format allowing immediate recovery in case of accidental loss
- New alert management
- Compatible with Apple Automator
- View and edit backup scripts
- Test Scripts before running them
- Backups can be launched when no session is open
- Intego calendar for Apple's iCal
- Improved backup statistics with graphs
- Backups can be made to multiple CDs or DVDs
- Automatic mount and dismount of network volumes



Apple Store

MacMall

NAVARRE

Fry's

MUSIC WORLD

www.intego.com

Intego • 500 North Capital of Texas Hwy, Suite 8-150 • Austin, TX 78746 • Tel (512) 637-0700 • Fax (512) 637-0701 • sales@intego.com

BlackBerry, Mac, the Mac logo and Apple Mail are registered trademarks of Apple Computer, registered in the U.S. and other countries. © 2007 - 2008 Intego, the Intego logo, Personal Backup, the Personal Backup logo and we protect your world are registered trademarks of Intego, registered in the U.S. and other countries. All other brand and product names are the property of their respective owners. Intego is not responsible for omissions or errors in typography or photography. Photo: Intego.



we protect your world

`link.title` is the name of the submit button for your text input box. `description` provides user instructions for the text input box. `name` is the form name of the text input box. `link` is the URL of your server-side script (such as PHP or Perl) that should process these text input submissions. See **Listing 1** for the XML code syntax of the `textInput` tag and its nested sub-elements.

## The Anatomy of RSS Feed Items

After you've defined your channel elements within the `channel` tag, it's now time to add your actual content items. These are the items that are displayed by aggregators as news stories, blog entries, podcast items, etc. (depending on what content you wish to include in your feed). Each item is encapsulated in its own `item` tag (that is nested within the `channel` tag below the channel elements). See **Listing 1** for the XML code syntax of the `item` tag and its nested sub-elements. The sub-elements that describe an item's content are defined here.

While most feed items include a link back to the full online version of the article or blog entry, a feed item does not require a link if you wish to include all of the content in its `description` sub-element. In fact, none of the `item` sub-elements are required as long as you include at least the item's `title` or `description`.

### title

Optional. This is the title of the item. Although it's optional, most aggregators look for this item sub-element, so it's highly recommended to include it.

### link

Optional. This is the URL to the web page version of the item on your web site or blog.

### description

Optional. This is the description of the item. For news stories and blog entries, it's your choice to include the entire text or only a summary. If you only include a summary, then be sure to also include the `link` sub-element in your item, so that users can click-through to your site to read the entire story. For podcasts, the `description` sub-element usually holds text information about the song track and artist. Although it's optional, most aggregators look for this item sub-element, so it's highly recommended to include it.

### enclosure

Optional. This sub-element defines a multimedia file. If your feed is a podcast, then the `enclosure` sub-element of each item would identify the related audio file. It requires three attributes: `url`, `length`, and `type`. The `url` attribute should be the direct URL to the actual media file itself. The value of `length` should be the file size of the media object in bytes. The value of `type` should refer to the media file's

MIME type. See **Listing 1** for the XML code syntax of the `enclosure` sub-element and its attributes.

### category

Optional. This sub-element works exactly like the channel's `category` tag, except that it defines a unique category for the individual item. Like the channel's `category` tag, there is no standard cataloging system, so often this tag only proves useful for your own site needs.

### source

Optional. This should name the source of the item's content if you are not the original author. For example, if the content is from MacTech's RSS feed, you would list "MacTech" as the `source` value and the URL for MacTech's RSS feed as the value of the `url` attribute.

### author

Optional. The e-mail address of the author of the item's content. For example, if the `source` is credited to MacTech's RSS feed and David Sobsey is the author of the piece, then list David Sobsey's e-mail address as the value of this sub-element.

### pubDate

Optional. This is the publication date of the item. Like the channel's `pubDate` and `lastBuildDate` tags, this sub-element's date should be formatted to conform to RFC 822.

### comments

Optional. This is the URL for the item's related web page of user comments. This sub-element is usually only relevant for blogs that allow web-based user comments.

### guid

Optional. This is an interesting sub-element that is used as a unique identifier for the item. There is no set convention for how this sub-element should be used, but most aggregators expect it to be a unique URL string that no other item can have, making it a valid item ID. For news stories and blog entries, this works great since they would have their own unique URLs, but if the item's URL is referred to more than once in your feed, then it cannot be used as the unique identifier here. For a unique URL that will always be available for viewing online, you should include the `isPermaLink="true"` attribute.

## Moving Forward

Now that you've stepped through the entire RSS 2.0 specification, you're ready to put RSS to good use. If you're only interested in building your own RSS feeds, then check out the Resources section of this article for links to Mac-compatible feed generator applications that can help streamline the creation process. If you're comfortable working with XML, you can also use a standard text editor

QUANTUM. WE'VE GOT YOU COVERED.



**The PX500 Series works weekends and holidays, so you don't have to.**

With shrinking backup, recovery and archive windows, most IT Professionals protect their data after normal business hours and on weekends – the times when you would rather not sit around watching a tape library. Our PX500 Series redefines value in rackmount tape automation with enterprise-class features, high density and market-leading investment protection. And our superior performance, reliability and support allow you to have a normal life – with vacations. To find out how Quantum's got you covered with our new PX500 Series, call 866-827-1500 or visit us at [www.quantum.com](http://www.quantum.com).

**Quantum.**



**BACKUP. RECOVERY. ARCHIVE. IT'S WHAT WE DO.™**

©2006 Quantum is a trademark of Quantum Corporation in the United States and other countries. All other trademarks are the property of their respective companies.

like BBEdit (<http://www.barebones.com/>) or an XML editor like <Oxygen/> (<http://www.oxygenxml.com/>) to roll your own RSS code. To ensure that your feed adheres to the RSS specification and does not include any errors, always test your feed with one of the many online validators (see the Resources section for some helpful URLs).

If you're interested in parsing RSS in your own software projects or syndicating third-party feeds on your web site, the Resources section also includes links to a few of the most popular RSS parsers for various programming languages. When developing your own RSS-savvy application or web site, it's important to assume that all RSS feeds are invalid until proven otherwise. Because many of the feeds out there are hand-coded, they often contain the wrong text encoding attribute in the xml tag and/or just bad XML, so it's a good idea to include a lot of error handling in your own parsing code to safeguard your users from problems. Defensive programming is definitely the name of the game here. It's also recommended to support caching in your RSS applications. If third-parties are kind enough to offer you free content for syndication, be considerate of their server bandwidth by caching the feed content temporarily on your end and only retrieving a fresh feed at timed intervals.

For those of you who want more out of RSS, the 2.0 specification allows you to extend RSS with namespace-defined modules. If you're not familiar with namespaces, they are part of the XML specification that RSS 2.0 supports as an easy way to extend the feed format for custom purposes. For more information on using namespaces to extend RSS 2.0, visit: <http://www.reallysimplesyndication.com/howToExtendRss>

And last, but not least... once you have your own RSS feed available on your site, it's easy enough to add a link to it on your home page using one of the popular "chicklet" buttons shown in **Figure 1**, but how do you get Safari and Firefox to automatically recognize your feed (see **Figure 3**) with that dynamic feed icon in the location bar and status bar respectively?



**Figure 3. Dynamic feed icons in Safari and Firefox.**

The answer is actually quite simple, requiring only a single line of HTML code. Nested within the head tag of your web pages, add the following meta tag with the href attribute set as the direct URL to your RSS feed:

```
<link rel="alternate" type="application/rss+xml"
title="RSS" href="http://yoursite.com/feed.rss">
```

So what are you waiting for? Now that you've gained some insight on how RSS 2.0 works, start taking advantage of one of the hottest new Internet technologies that can greatly

increase exposure for your products and services, and drive additional traffic to your web site or blog.

## Resources

The following list is by no means comprehensive, but should serve as a good starting point for learning more about RSS 2.0 online.

### RSS Readers/Browsers for Mac OS X

NetNewsWire: <http://ranchero.com/netnewswire/>  
Radio Userland: <http://radio.userland.com/>  
NewsFire: <http://www.newsfire.com/>  
Apple Safari: <http://www.apple.com/macosx/features/safari/>  
Firefox: <http://www.getfirefox.com/>

### RSS Feed Generators

FeedForAll: <http://www.feedforall.com/>  
ListGarden: <http://softwaregarden.com/products/listgarden/>  
pReSS: <http://www.mizog.com/productinfo/press/>  
OrangeBox: <http://www.globalsyndication.com/orangebox-for-macintosh>  
Feeder: <http://www.reinventedsoftware.com/feeder/>

### Online RSS Validators

Userland Validator: <http://rss.scripting.com/>  
Feed Validator: <http://feedvalidator.org/>

### RSS Parsers

MagpieRSS for PHP: <http://magpierss.sourceforge.net/>  
CaRP for PHP: <http://www.geckotribe.com/rss/carp/>  
Universal Feed Parser for Python: <http://feedparser.org/>  
RSS Class for Cocoa: <http://ranchero.com/cocoa/rss/>  
WSL-Feed for REALbasic:  
<http://www.ebutterfly.com/rb/webservices.asp>

## Documentation

RSS 2.0 Specification: <http://blogs.law.harvard.edu/tech/rss>

## Books

*Developing Feeds with RSS and Atom* by Ben Hammersley (O'Reilly):  
<http://www.oreilly.com/catalog/developrssatom/>  
*Beginning RSS and Atom Programming* by Danny Ayers and Andrew Watt (Wrox):  
<http://www.wrox.com/remtitle.cgi?isbn=0764579169>



## About The Author

*Dave Wooldridge is the founder of Electric Butterfly ([www.ebutterfly.com](http://www.ebutterfly.com)), the developer of the Web Services Library for REALbasic and the award-winning HelpLogic. He is also co-founder of the new eBook publisher, SpiderWorks ([www.spiderworks.com](http://www.spiderworks.com)).*

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# NAGIOS ON OS X, PART 1

## INSTALLING AND SETTING UP ONE OF THE BEST NET MONITORING TOOLS AROUND

If you read my website, <http://www.bynkii.com/>, you may know I have talked about installing Nagios in the past. For a while now I've wanted to write an updated, in-depth piece on this subject, and felt now was time to do so, and of course, only in MacTech! I've managed to simplify the install quite a bit, thanks to the hard work of the folks at DarwinPorts, <http://darwinports.org>.

### Welcome

This article is going to have a few sections. First, we'll get DarwinPorts installed and configured. Then we'll use it to install the backend components Nagios needs. Third, we'll do the actual Nagios installation, and the plugin installation. Then we'll go over some of the basics of configuring Nagios. Note that this applies to the current version, which is 1.2. Version 2.0 is in final beta, but I prefer to wait until that's done before installing it on production systems.

We should also talk a bit about what Nagios is. Nagios is a network monitoring tool. It allows you to monitor services running on various hosts on your network. For example, you can monitor your switches to ensure they're working, or you can monitor various critical server processes on an Xserve, like the KDC process, AFP processes, etc. You can, with various plugins from the Nagios Exchange site, at <http://www.nagiosexchange.org/> monitor internal counters on Windows servers too. If you're skilled with Perl, or any one of dozens of programming languages, you can write your own plugins.

### Darwinports

DarwinPorts, like Fink, or any other of a dozen ports management systems is a way to

make installing and configuring software easier. If you compile and install open source software manually, you have to figure out all the various dependencies and configuration issues yourself. This is not hard per se, but it is tedious. Ports managers, like DarwinPorts manage this for you. If the DarwinPorts repository has the software you want, you install it, and it handles all the dependencies for you. While, like everything else, there is any number of religious wars about ports systems, I like DarwinPorts, because it does the job well for me. If you like Fink better, great! They're both really good systems.

To avoid modifying Apple-supplied directories, DarwinPorts lives in `/opt/local/`. If you install on Mac OS X Server, then you'll already have `/opt`: it's where Apache2 is installed. The advantage to this is that it makes uninstalling DarwinPorts dead simple. Remove `/opt/local` and all its contents, and DarwinPorts is gone.

Since DarwinPorts doesn't have a Nagios port, we can't just use it for the full install, but we can use it, and some other tools for the support libraries and applications Nagios needs. However, the first thing we have to do is install DarwinPorts itself. The easiest way to do this is to get the disk image with the installer from DarwinPorts, currently at <http://darwinports.org/downloads/DarwinPorts-1.1.dmg>. Download

# Windows XP on a Mac? *Maybe pigs can fly!*

## Apple **Boot Camp** Public Beta **First Look**

Get all the info about Boot Camp, first.



Apple Boot Camp Public Beta First Look eBook can help your Mac Fly too!

### You will learn how to:

- Get Boot Camp up and running on your Intel-based Mac
- Avoid Windows viruses and malware
- Configure your Windows system
- Share data with your Mac partition
- Back up your Windows system
- Run a virtual copy of Windows within the Mac OS
- Much, much more...

All for only **\$6.99**, simply visit  
**[www.peachpit.com/bootcamp](http://www.peachpit.com/bootcamp)**  
to download your copy today!



**[www.peachpit.com](http://www.peachpit.com)**

the disk image, and run the installer. Once you're done with that, we're going to be doing the rest of our work in Terminal.

To make life simpler, we're going to set up a .profile in our home directory that will make using DarwinPorts easier. While you don't have to do this, if you have multiple versions of software installed in various places, a .profile file will make your life a lot simpler. My .profile is a bit simplistic, but it works for my needs. In any text editor that will allow you to create .files, (pico, vi, emacs, TextWrangler, SubEthaEdit, BBEdit, they're all good) create a file with the following entries, and save it in the root of your home directory as .profile:

```
PATH="/bin:/sbin:/usr/bin:/usr/sbin:/usr/local/share
/aclocal:/opt/local:/opt/local/bin:/opt/local/lib:/o
pt/local/sbin"
```

```
export PATH
```

```
PKG_CONFIG_PATH="/opt/local/lib/pkgconfig"
```

```
export PKG_CONFIG_PATH
```

```
MANPATH="/usr/share/man/:/usr/X11R6/man:/opt/local/
share/man"
export MANPATH
```

```
INFOPATH="/opt/local/share/info"
```

The path statement ensures that when you enter a command, that the /opt/local tree is used automatically. It also leaves the standard paths for /bin, /sbin, /usr/bin, /usr/sbin, etc. The PKG\_CONFIG\_PATH is used to ensure that software you install under your login can find pkg-config, a tool used to make compiling libraries and applications easier. The MANPATH statement makes finding man pages in the /opt/local/ tree easier. The INFOPATH variable is used by things like gettext, which figures into a lot of open source packages. If you already have your own .profile, then just add the /opt/local information if you wish. I rely heavily on DarwinPorts, so for me, having this .profile setup is a real timesaver. If you don't use DarwinPorts a lot, you may not wish to modify your .profile.

So, once you're done setting up, (or not) your .profile, we're ready to go. If you did modify your .profile from within the Terminal, you'll want to read that file so your environment is set up. Just run

```
source ~/.profile
```

to set up your environment with the new information from .profile. Now we need to make sure DarwinPorts is up to date. To do this, run:

```
sudo port selfupdate
```

and wait a few minutes. If there is any updating to do, you'll see some short status messages on the screen, and if all goes well, you'll get the "selfupdate done!" message, and you're set. If you want to see *all* the status messages for a selfupdate, then run:

```
sudo port -d selfupdate
```

which enables debug mode for that command, you'll get all the status messages. Below, I have the two screenshots for the same command on an updated system to give you an idea of the difference the -d makes.



Figure 1: Normal selfupdate




Figure 2: Debug selfupdate

I tend to only use the -d option if I'm having problems, so I can better see where the problem is occurring. I like to run the selfupdate option at the beginning of any install session, so I know I'm starting with the most current version of DarwinPorts.

Once that's done, we install the various background libraries that Nagios needs using

# Introducing SyncDeK® 6



## The Premier Data Replication Engine for FileMaker® *Just Got A Whole Lot Better!*

### IDEAL FOR:

- Continuous Server Backup and Recovery
- Warm Standby
- Server Clustering
- Multiple Offices & Remote Sites
- Mobile Users
- Web Server Sync
- Distributed Versioning

### Data Protection:

Protect mission critical data with continuous differential data backup to one or more offsite locations. Setup Warm Standby and fail-over servers to ensure access to users at any time.

### Mobility:

Deliver fresh data to users anytime, anywhere in the World. Whether your users work on the go, log in from home or work in multiple locations, SyncDeK gives each location its own local copy of the database to access, whether users are on or offline.

### Security:

SyncDeK encrypts all data transmissions and allows you to place your actual databases inside of firewalls to prevent intrusion.

### What's New In SyncDeK 6:

- Continuous Background Sync on servers and laptops - No user interruption
- Direct integration with FileMaker - No scripting or import/export setup
- SyncDeK Server 6 - Centralized, Secure, Scalable Data Replication
- FileMaker Server Advanced Replication - No client required
- Automatic Container field replication
- Version Update Manager for remote FileMaker Servers



### FILEMAKER DEVELOPERS

Visit our website to find out more about how SyncDeK benefits organizations that rely on FileMaker Pro

Learn about how the SyncDeK Developer Alliance can support your business

For more information, including: **Special pricing, demo software, and a FREE technical consultation, contact us NOW !**

**877-548-4920** (toll-free) U.S.

**510-548-4920** International

**WorldSync.com/SD6**

DarwinPorts and the install switch:

```
sudo port install zlib
sudo port install libpng
sudo port install jpeg
sudo port install gd2
```

Once all of those are done, congrats, you've installed the support libraries for Nagios. Note that there are *tons* of ports in DarwinPorts, so this is only the barest smidgen of what you can get with DarwinPorts.

We'll need to install at least one particular Perl package for Nagios, namely Net::SNMP. To do this, run:

```
sudo cpan -i Net::SNMP
```

If you've never run `cpan` before, you'll have to run through some configuration steps first. They're pretty easy to follow, and if you're unsure of what to do in a given step, take the defaults. There's even an option to let `cpan` try to autoconfigure itself. I do it manually, but if autoconfigure works for you, so much the better. Using `cpan` is much like using DarwinPorts, only just for Perl. If you're trying to install a Perl module that has dependencies that you don't have installed, `cpan` will prompt you to install those as well. (If you've only ever hear horror stories about downloading and installing open source software, I'm here to tell you that thanks to a lot of hard work by not a lot of people, that process has gotten much easier, and thanks to tools like `cpan` and DarwinPorts, should no longer be thought of as scary to the uninitiated.)

## Installing Nagios

Now that we have the support libraries installed, it's time to set up for Nagios. We'll need to create the directory that Nagios lives in, `/usr/local/Nagios`. Next we have to create the users and groups that Nagios will use. (Security note: *All* of this assumes that you have a proper firewall and other protective measures in place. I would *highly* recommend that you don't install Nagios on a system directly exposed to the public Internet. A lot of the protocols that Nagios uses, like SNMP v2c are not encrypted or terribly secure, and having a Nagios box exposed to the public Internet will potentially create real problems for you.)

You'll need to create at least a Nagios user and a Nagios group. Note that the nagios user does not need a login shell, just that the account works. Regardless of how you create the nagios user account, (System Preferences, NetInfo Manager, Workgroup Manager, command line), you want to make it a local machine account *only*, and you want to set the login shell to `/usr/bin/false`. Don't give the account a password. You'll also need to create a nagios group. The Nagios group needs two members, nagios and www. (Adding www to the nagios group makes setting

up the integration between Apache and Nagios much easier.) Once you've created the user and group for Nagios, then we're ready to download the Nagios source code files.

Go to the Nagios download site, <http://www.nagios.org/download/>, and download the 1.2 version of the Nagios tarball, and the most recent version of the Nagios plugin tarball, 1.4.2 as of this writing. Save them both to wherever you want, I usually save them to my desktop. Unzip and untar the Nagios source code files. In the Terminal, `cd` to the nagios-1.2 directory. The first thing we have to do is run the configure utility so that we can compile and build the code correctly. As we're using some non-standard (for Nagios) library locations, we'll need to tell it where to find things. We also need tell Nagios where its own base directory is going to be, where the CGI directory it will use is going to be, and where the base web root directory it's going to use is. For this article, I'm using:

```
./configure --with-gd-lib=/opt/local/lib --with-gd-inc=/opt/local/include --prefix=/usr/local/nagios --with-cgiurl=/cgi-bin --with-htmlurl=/
```

Make sure you don't have any spaces between the `=` and the leading `/` of the paths in the configure command, and let it rip. When you get done, you should not have seen any warning `s` or errors during the configure, and you should see a status screen like the one below:




Figure 3: Completed configure screen


You'll want to save this screen somewhere, in case you aren't sure where things are. With this setting, Nagios is going to use the defaults for things, so we'll have to change a few things. First, it assumes that you're going to have all your web files originating from `/nagios` in the URLs. We'll fix that later. The `cgi-bin` directory is the system default one, or `/Library/WebServer/CGI-Executables`. That we'll leave alone.

# YOUR ULTIMATE 3D AQUARIUM!



**FORTY BEAUTIFUL FISHES,  
TURTLES, JELLIES, & SHARKS!  
INTERACTIVE SCREENSAVER  
TAP THE GLASS, FEED THE FISH**

 **Allume  
Systems**  
A Division of Smith Micro Software



**A virtual aquarium so vibrant, you'll swear the fish are real! Available from your favorite retailers and catalogs.**

Aquazone Seven Seas Deluxe includes 40 fresh and saltwater fish, jellyfish, adult and baby sea turtles, and sharks! Have entire schools of fish on screen swimming through 20 customizable 3D tanks!

® TM, and © 2005 Allume Systems, Inc. To learn more about our range of software solutions for your Macintosh, please visit us at [www.allume.com](http://www.allume.com).

Assuming your configure went okay, we're going to run the various Nagios make and install commands:

```
make all
```

If make all completes successfully, you'll see a status screen like the one below. We're going to run the other make commands with the exception of make install-init, since Mac OS X doesn't handle startup items like that, and we'll want to test a bit before we make Nagios a startup item.



Figure 4: Completed make all screen

```
sudo make install
```

If the make install goes well, you'll get another status screen telling you so.

```
sudo make install -commandmode
```

This command sets up the external command directory for you, important if you want to extend Nagios beyond the basics

```
sudo make install-config
```

This sets up the sample configuration files, which we'll go into later. Next is to configure and install the plugins. As with Nagios, unzip and untar the source files for the Nagios plugins. Change directory into the nagios-plugins directory and run `./configure`. If there are no errors, run `make`, then `sudo make install`.

## Initial Post – Install Setup

So Nagios is installed, the plugins are installed, we just start and go, right? Well, not so fast. First, we have to make sure that the web server side of things is set, since Nagios uses a web interface to show you what its monitoring. If you installed Nagios on Mac OS X Server, then you just use Server Admin to point the web server root at `/usr/local/nagios/share`. You'll want to make sure you have CGI Execution enabled for the site, and that you have the necessary Apache modules enabled to allow you to run Perl and other language CGIs on the site. If you are using Mac OS X instead of Mac OS X Server, you'll have to edit the necessary Apache config files by hand to do this. Oh yes, make sure that evil Performance Cache is turned off.

Next, we copy the Nagios CGIs into `/Library/WebServer/CGI-Executables/`. The CGIs can be found in `/usr/local/nagios/sbin`. You want to make sure that at least the group for the CGIs is set to `www`, and that the permissions are set so that the owner and group can execute. (*I'll leave world executable rights up to your particular needs and security posture*).

The next step is to ensure that Nagios knows where its various web files really are. To do this, we want to open up `side.html` in `/usr/local/nagios/share`, and make sure that all the references to CGIs look like this: `href="/cgi-bin/statuswrl.cgi`. If they do, then you're all set there.

If you've got everything pointed right, then we can test the basic operation, aka "can we see the Nagios Home Page". Just point your browser at the Nagios URL, and if it's all set up correctly, you'll see:




Figure 5: The Nagios Home Page

Note that nothing works yet, and in fact, the Nagios process isn't even running. But at least we


# The four D's of 4D



Design truly beautiful, native applications with rich, intuitive user interfaces that users demand to maximize their productivity.



With a powerful and efficient development environment, you can create powerful apps that take advantage of all of today's standards.



With an integrated compiler, automatic backup and client updating, support, delivery, and installation of your apps has never been easier!



Go ahead. Kick your feet up. You can have peace of mind knowing that you're getting the best in performance and ease-of-use.

## 4th Dimension 2004

[www.4D.com](http://www.4D.com)

know it can see the home page. So that's something. Next we want to make sure that just anyone cannot see all the information that Nagios can provide. If nothing else, you don't want the entire world, or even your entire company, getting detailed information on your servers.

There are a lot of ways to set up access control to web sites. I'm covering one here, namely using .htaccess files. To do this, we're going to make a couple of changes to httpd.conf, and set up a .htaccess file for locking down access.

httpd.conf changes:

*Note: This is a fairly simple way to set this up. There are other ways to do the same thing, and they will work just as well. This one works in my situation.*

?Open /etc/httpd/httpd.conf in your text editor of choice (I use [BBEdit](#)) and look for the following directive:

```
<Directory "/Library/WebServer/CGI-Executables">
  AllowOverride None?
  Options None?
  Options ExecCGI
  Order allow,deny
  Allow from all
</Directory>
and change it to read:
<Directory "/Library/WebServer/CGI-Executables
#AllowOverride None
AllowOverride AuthConfig
# Options None
Options ExecCGI
Order allow,deny
Allow from all
</Directory>
```

This will allow you to run the CGI's after you have authenticated yourself for the CGI's. To do that, we need to set up a .htaccess file in the cgi-bin directory, and then set up the passwords. Create a file named .htaccess in /Library/WebServer/CGI-Executables/ and set it up thusly:

```
AuthName "Nagios Access"?
AuthType Basic?
AuthUserFile /usr/local/nagios/etc/htpasswd.users
?require valid-user
```

Once you've done that save it and close it. Now, before we create the password file for this, you want to know all the users that will be authenticating. At the very least, you want some sort of catchall Nagios admin user. For our example, we used nagiosadmin, jwelch, and admin

Run the following:

```
sudo htpasswd -c
/usr/local/nagios/etc/htpasswd.users nagiosadmin
```

This will create the user file that our .htaccess file is looking for, and set the first user to nagiosadmin. You'll be asked for the password for this user, set it to what you feel is correct.

# macforge.net™

MacForge indexes and tracks open source projects that run on the Mac, or are likely to without modification. Thanks to MacForge, there's no need to sift through huge listings of open source that you can't use. With categories, filters, and more, MacForge makes it easy to find what you need.



**MacForge:**  
Your Gateway to Mac Open Source  
[www.macforge.net](http://www.macforge.net)

Sponsored by **MACTECH**



# BookEndz

**17"  
Now  
Shipping!**

## DOCKING STATIONS

**MacBook  
Pro Dock  
available in  
July!**

Convert your PowerBook or iBook to a desktop system in seconds without misplacing cables or damaging connectors.

New Higher Resolution  
G4 15" - 17" PowerBook  
Docking Station



- Eliminates cable confusion and damage to connectors.
- New release levers for easier docking and undocking.
- Adds only 3" to rear when docked.
- Also available for G3 PowerBooks.

iBookEndz  
Docking Station



12" G4 PowerBook New Higher Resolution Docking Station

- Aluminum look to match your PowerBook.
- Eliminates cable confusion and damage to connectors.
- All connectors are routed to the rear of the Dock.
- The Dock converts RGB to standard VGA connector.
- Streamlined design complements your PowerBook.
- Use internal or external speakers.
- Easy to use ejection system.



Check our Web Site for latest product announcements



**BookEndz – Manufactured by OlympicControls Corp.**

**1250 Crispin Drive, Elgin, Illinois 60123**

**Phone: 888-622-1199 • Fax: 847-742-5686 • [www.bookendzdocks.com](http://www.bookendzdocks.com)**

**[www.bookendzdocks.com](http://www.bookendzdocks.com)**

To add other users, you run:

```
sudo htpasswd /usr/local/nagios/etc/htpasswd.users  
<username>
```

You do this once for each user you want to add. You only need the -c switch when you are creating the file and first entry.

Then, edit the "use\_authentication" line in /usr/local/nagios/etc/cgi.cfg (or cgi.cfg-sample) to read:

```
use_authentication=1
```

If you read that file, you'll also note this is where you set which nagios user can do what. The docs that are a part of Nagios explain this thoroughly, and are available from the "Documentation" link on your Nagios home page, so I don't have to.

## Conclusion

I want to stop the article here, because the next part is going to go into the configuration files, which covers a *lot* of detail, and can be pretty dry. Unfortunately, there's no nice or easy way to jump in to that. I *highly* encourage anyone who wants to go ahead on their own to do so, but with caution, the config files can get a little odd at times. I'd also encourage any readers who are not going to plow ahead to read the Nagios documentation thoroughly. Nagios has its own language, and the more familiar you are with it, the happier you'll be using it. As always, thanks for reading!

## Bibliography and References

DarwinPorts, <http://darwinports.org>.

CPAN, the Comprehensive Perl Archive Network, <http://www.cpan.org/>

Nagios, <http://www.nagios.org/>

Nagios Exchange, the Nagios Plugins and Add Ons Exchange, <http://www.nagiosexchange.org/>



### About The Author



**John Welch** <jwelch@bynkii.com> is Unix/Open Systems administrator for Kansas City Life Insurance, (<http://www.kclife.com/>) a Technical Strategist for Provar, (<http://www.provar.com/>) and the "GeekSpeak" segment producer for Your Mac Life, (<http://www.yourmaclife.com/>). He has over fifteen years of experience at making Macs work with other computer systems. John specializes in figuring out ways in which to make the Mac do what nobody thinks it can, showing that the Mac is a superior administrative platform, and teaching others how to use it in interesting, if sometimes frightening ways. He also does things that don't involve computertry on occasion, or at least that's the rumor.



# Speech Recognition for Mac OS X

## New Release! iListen 1.7

- New! Universal Binary**
- New! Faster Training**
- New! Improved Voice Commands**

### Power of Speech

MacSpeech's exclusive TalkAnywhere™ technology allows you to use your voice to enter text into virtually any application.

### Faster Training

With the improved training and set-up process in iListen 1.7, you can now be dictating in as little as 5 minutes.

### Always Learning

CorrectAnywhere™ - another MacSpeech exclusive lets you use Correction without leaving the application you are using. This allows iListen to learn new words and improve your voice profile while you use the program.

<http://www.macspeech.com>



**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# Build Your Own Supercomputer From Your Macs Laying Around

## How to employ Tiger's Xgrid to build clusters and to contribute to grid/cluster projects.

By Mary Norbury-Glaser

### Grid vs. Cluster

Even in the Linux world, where cluster and grid computing are well established, there is often confusion about how to distinguish a grid from a cluster. Let's look at characteristics of each.

Grid computing uses the available resources of many individual, loosely coupled computers used by many different people across organizations. Grid computing makes use of idle computer time and unused disk space on different systems, often desktop computers (the SETI@home project, for example). Administrative tasks like system scheduling and job management are "distributed" since the computers are typically located over a wide area of multiple domains. There is usually no single system image; Windows, Linux and Mac systems often contribute to a single grid. Grid computing can be used for both high throughput computing (Apple/Genentech BLAST, for example) and high performance computing (the San Diego Supercomputer Center, the National Center for Supercomputing Applications, the Argonne National Laboratory, and the Max Planck Institute for Gravitational Physics collaborated in the largest grid computing demonstration of simulations involving Einstein's General Relativity equations).

Cluster computing, in contrast, involves a group of tightly coupled computers that work in parallel to share processing as if they were one machine with multiple CPUs. Administrative job management and a scheduling system are centralized in a cluster. Cluster machines also share a single system image. In other words, the collection of systems in a cluster appear as a single entity to the user, the DBA, etc. Cluster computing can be used for high performance computing (see description in grid computing above), load balancing and high availability computing. Load balancing describes the situation when processing activity is efficiently distributed among cluster members so no individual computer gets overwhelmed. High availability computing is implemented in situations that require maximum uptime and where availability of services is critical.

Think of grid computing as involving computers from myriad networks around the globe while cluster computing is a set of machines usually in one location involved in parallel

computing with the goal to get them to look like a single virtual machine.

### Xgrid Introduced

At MacWorld 2004 in San Francisco, Apple's Advanced Computation Group quietly announced the Technology Preview Release of Xgrid version 1.0 as a free beta download. Xgrid created quite a lot of excitement in ripples throughout the conference attendees who were involved in or interested in compute intensive applications. The software download included the Xgrid app and a basic local alignment search tool: Apple/Genentech BLAST (an open source DNA and protein sequence matching application) that enabled distributed searches on an Xgrid cluster. It touted the use of zero configuration (Rendezvous at the time, now Bonjour) to discover available resources on the network.

XGrid is based on a NeXT application called Zilla developed in the late 80's by Richard Crandall, a then NeXT employee and now a Distinguished Scientist at Apple. Zilla was the first community supercomputing application (screen saver type distributed computation system).

With the publicity of the Virginia Tech's Terascale Computing Facility and their System X built on 1100 dual processor, 2.0 GHz Power Mac G5 computers, the potential of low-cost supercomputing became an exciting possibility. Using commodity based or commercial off-the-shelf (COTS) hardware and free software, distributed computing became well within the reach of any organization with existing hardware.


### Xgrid and Tiger

In April 2005, Apple introduced Mac OS X version 10.4 codenamed "Tiger" and included Xgrid in both the client and server versions of the operating system. Xgrid was streamlined and arrived in Tiger a different animal:

Xgrid in Tiger no longer limited job submission to the GUI; the client-side Cocoa API was made available to developers who were encouraged to assimilate Xgrid into their applications

# Complete Source Control and Defect Management for Mac OS X

**Seapine Software™**  
changing the world  
of software development



## Effective source code control and defect tracking require powerful, flexible, and easy-to-use tools—Surround SCM and TestTrack Pro

- Complete source code control with private workspaces, automatic merging, role-based security, and more
- Comprehensive defect management — track bug reports and change requests, define workflow, customize fields
- **New!** Full Unicode support maintains international characters across various languages, ensuring data is not lost or misinterpreted
- Advanced branching simplifies managing multiple versions of your products
- Fast and secure remote access to your source files and defects — work from anywhere
- Scalable and reliable cross-platform, client/server solutions support Mac OS X, Windows, Linux, and Solaris
- Exchange data using XML and ODBC, extend and automate with SOAP support
- Licenses priced to fit your budget

**Seapine Software Product Lifecycle Management**  
Award winning, easy-to-use software development tools

Seapine  
**Surround SCM**  
Seapine  
**TestTrack PRO**



Download Surround SCM  
and TestTrack Pro at  
[www.seapine.com/mac](http://www.seapine.com/mac)  
or call 1-888-683-6456

all product names listed herein are registered trademarks of their respective owners. All rights reserved.



instead of writing plugins. There are already several clients that have emerged from this addition to Xgrid: GridStuffer (a Cocoa client available at <http://cmgm.stanford.edu/~cparnot/xgrid-stanford/html/goodies/GridStuffer-info.html>) and PyXG (a Python interface to Xgrid, <http://hammonds.scu.edu/~classes/pyxg.html>).

A frustrating change for many users of the Panther technology release of Xgrid is the loss of the Xgrid **controller** on the Tiger client. Tiger client still includes the Xgrid **agent** but the GUI **controller** has been moved to Tiger server (see the next section for definitions of Xgrid terms). However, when we build our own cluster, we'll have some options. Patience, dear reader...

## Definitions and Simple Explanations

**Client:** the Client submits jobs to the Controller. This can be any OS X 10.4 computer or OS X Server.

**Controller:** the Controller receives the jobs from the Client(s), splits the jobs into tasks and submits the tasks to the Agent(s). The Controller then receives the results back from the Agent(s) and delivers the information back to the Client(s). Apple has moved the Controller to OS X Server but has left the command line equivalent for managing the Controller and job submission on Client.

**Agent:** the Agent receives tasks from the Controller, runs the computations and sends the results back to the Controller. The Agent can be either a 10.3 or 10.4 computer. One task per CPU can be run.

**Node:** any single OS X (10.3 or 10.4) desktop computer or OS X 10.4 Server on the network.

There is only one Controller but there can be many Clients and many Agents.

The Controller, Client and Agent can be run on the same machine for purposing of testing. This is not, however, an optimal scenario for real world application.

## Participate in a Grid...

Grid computing allows for an environment where idle CPU cycles and storage space of thousands of networked systems can work together on a particular processing-intensive problem. Current projects range from mathematics (prime number searches), science (climate prediction models), life sciences (cancer research), to cryptography (cracking data encryption schemes). By joining a grid project, your home or work computer can contribute to solving a global challenge.

The simplest use of Xgrid is to add your Mac to an existing distributed computing grid project. Deciding to participate in a grid is as simple as looking online for available Xgrid projects: <http://distributedcomputing.info/projects.html> lists active and upcoming projects under different categories with details on supported operating systems.

For example, Rosetta@home (<http://boinc.bakerlab.org/rosetta/>) is a project that is attempting to predict and design protein structures in an effort to help cure human diseases. To join this grid, go to their home page, create an account (you'll receive a return Account Key via email which you'll need to confirm at the web site and later insert into the application), make sure your computer meets the system requirements, and download the BOINC client software (Berkeley Open Infrastructure for Network Computing, <http://boinc.berkeley.edu/>) in either the GUI or command line flavor. The client software enables computers with different operating systems to "talk" to the project server.

Select the Projects tab and click the Attach to new project button. You'll be prompted for the project URL and Account Key that you received when you created your BOINC account.

Based on your chosen preferences (processor, disk, memory and network usage, etc.), the project will start and you can keep track of the progress through the tabs at the top of the BOINC app window.

Instructions for the command line version, along with command line options are available here: <http://boinc.berkeley.edu/download.php>.

## ...Or a Cluster

In this tutorial, we'll be looking at Xgrid@Stanford, a project run by Charles Parnot, a postdoctoral fellow in the Molecular and Cellular Physiology Department at Stanford University. His research involves 3D modeling of G protein-coupled receptors to study heart disease and heart rate control. The Stanford project began in March 2004 with 8 computers running at about 4 GHz. By September 2005, they had nearly 500 registered agents with an average of 200 machines online at any given time, the cluster running over 200 GHz. Xgrid@Stanford's home page (<http://cmgm.stanford.edu/~cparnot/xgrid-stanford/index.html>) describes the project and details FAQs, Goodies and Latest News about the project. Note that the Stanford group calls their project a "cluster" because they are using Xgrid exclusively and all the participating agents are Mac OS X computers. To join Xgrid@Stanford, we'll start by introducing ourselves to Xgrid (note the team has provided separate instructions for Panther and Tiger):

### Set up Xgrid on Tiger:

1. Open System Preferences.
2. Open the Sharing Pref Pane.
3. Make sure your computer has a unique name!
4. Highlight Xgrid from the list under the Services tab and click the Configure button. Do not check the Xgrid checkbox!
5. Under Controller, select Use a specific controller: and type in:

**b161-g4.stanford.edu**

6. Under Agent accepts tasks: choose Always.
7. Under Authentication method: select None in the pop down list. You may get a security warning dialog box that you may ignore. Press OK on the Xgrid configuration sheet.



**Figure 1: Configuring Xgrid Pref Pane**

8. You will now be back at the Sharing Pref pane. Press the Start button.
9. Once the service is started, the Start button becomes a Stop button. If you need to change settings later, you will need to stop the agent, modify the settings and start it again.
10. Open the Energy Saver Pref pane and set Put the computer to sleep when it is inactive for: to Never.

Send an email to [charles.parnot@gmail.com](mailto:charles.parnot@gmail.com) with your computer name, where you are from and where you heard about the project (reference MacTech, please!). The project team will let you know about system updates.

#### Download Dashboard Widget:

Download the Xgrid@Stanford Tachometer Widget from the Dashboard Widget section of Apple's website (<http://www.apple.com/downloads/dashboard/status/xgridstanfordwidget.html>) or from the developer's site (<http://www.mekentosj.com/widgets/xgrid/>).



**Figure 2: Xgrid@Stanford Widget**

The widget shows the total number of active or inactive agents or processors, the percentage of working agents and the current cluster speed.

#### Restore Screen Saver Functionality:

In the developer's preview of Xgrid, there was a nice tachometer screen saver feature that disappeared in Tiger. To get it back and view the Xgrid@Stanford widget as a screen saver, download either of these screen saver modules that allow Dashboard widgets to move about your screen during idle time: **Amnesty** ([http://www.mesadynamics.com/amnesty\\_saver.htm](http://www.mesadynamics.com/amnesty_saver.htm)) or **Dashsaver** (<http://highearthorbit.com/software#DashSaver>).

In this example, I've chosen to download **Amnesty**:

1. Drop the Amnesty Screen Saver.saver file into your Screen Savers directory (~/.Library/Screen Savers).
2. Open the Desktop and Screen Saver Pref pane.
3. Select Amnesty Screen Saver from the list and click the Options button.
4. Select Xgrid@Stanford from the Widgets pop down list.
5. Set the time for when you want your screen saver to start.

## Roll Your Own Cluster

You only need two computers to build a cluster with Xgrid: one agent/controller and one other agent. In this tutorial, we'll use three computers: one 12" Powerbook (named tiger12), one Mac mini (named minimini), and one 15" Powerbook (our controller named norbury15).

**8 issues of MacTech Magazine for  
only \$29.95, over 60% off! Just visit:  
[store.mactech.com/sampler](http://store.mactech.com/sampler)**

## Network the Macs Together:

If your Macs are not already networked together, you can easily do so using a 10/100Base-T Ethernet hub and a cat-5 Ethernet patch cable for each Mac. Connect each Mac to a hub port. Your Macs are also networked if you are sharing an Airport connection but you'll obviously suffer speed degradation in this scenario.

Verify that Bonjour is enabled (in /Applications/Utilities/Directory Access.app) on all Macs.




Figure 3: Bonjour in Directory Access

## Configure the Agents (all Macs):

Open the Sharing Pref pane, highlight Xgrid from the list under the Services tab and click the Configure button. Select these options: Use first available controller, Always, none for authentication.



Figure 4: Configure the Agent

Click the OK button. You'll get a security warning; ignore it for now.

Click the Start button to turn on Xgrid Sharing.

## Configure the Controller:

Remember I mentioned previously that the GUI controller has been moved to Tiger server? We have two options for restoring this functionality in Tiger client: we can download **XgridLite** (<http://edbaskerville.com/software/xgridlite/>), a \$15 shareware add-on to Tiger Client's System Preferences or we can use the Terminal to control starting and stopping the controller.

**XgridLite** has some nice, basic features: it can turn the controller on and off and you can set passwords for client and agent authentication.



Figure 5: XgridLite

But since this is MacTech after all, let's do it from the Terminal using the xgridctl daemon. Xgridctl syntax comes in this flavor:

```
xgridctl status target on|off|start|stop|restart
```

where the target can be either c, which indicates the controller, or a, to indicate the agent. On/off refers to launching the daemon at startup and start/stop/restart explicitly controls the daemon.

Edit the file /Library/Preferences/com.apple.xgrid.controller.plist by changing the Agent Authentication string and the ClientAuthentication string from Password to None. Don't forget to use sudo!

```
sudo /usr/sbin/xgridctl c on
```

modifies the launchd configuration to run the controller daemon at startup and this command will start the controller:

```
sudo /usr/sbin/xgridctl c start
```

To use password authentication, you'll have to do a few more tweaks:

1. Edit the file /Library/Preferences/com.apple.xgrid.controller.plist by changing the Agent

***'Poof! Three months of work gone because I hit the wrong key.'***

Just because it's gone,  
doesn't mean it's gone for  
good. Get the hardest working  
data recovery software for  
Mac OS X on the job.

FileSalvage™ can undelete files  
from your hard drive, digital  
camera or iPod.

Even if you've accidentally  
formatted it or your drive is  
corrupted ... you can still  
get your important files back.

**See for yourself**

Visit [www.SubRosaSoft.com](http://www.SubRosaSoft.com)  
for more information.



**FILESALVAGE™**  
**SUBROSAFT.COM**

Authentication string and the Client Authentication string to use Password.

2. Go to the Sharing Pref pane for Xgrid and select the Password option for Authentication method: and enter the password you wish to use.
3. Reset the agent via the GUI (stop/start). This creates the file: /etc/xgrid/agent/controller-password
4. Copy the file:

```
sudo cp /etc/xgrid/agent/controller-password
/etc/xgrid/controller/agent-password
sudo cp /etc/xgrid/agent/controller-password
/etc/xgrid/controller/client-password
```

5. Start and stop both the agent (via the GUI) and the controller (via xgridctl as above).

The Xgrid Admin.app allows GUI management and monitoring of the controller and it's part of the Server Admin Tools. Server Admin Tools 10.4 can be downloaded and installed on a Tiger client machine (<http://www.apple.com/downloads/macosx/apple/serveradmintools104.html>). Put this on your designated controller (mine is norbury15).

Launch Xgrid Admin.app and you will see a sheet asking to enter or choose a controller. You should see the name of the computer that you started the controller on using xgridctl (norbury15, in my case).

In The Xgrid Admin window Click the Connect button and you'll see the Overview:

In the Overview window Click the Connect button and you'll see your agents listed under the Agents tab:




**Figure 6: Xgrid Admin Agents Tab**

There are no jobs listed under the Jobs tab yet so let's add one!

#### Add A Job Through the xgrid CLI Client:

Let's move over to my 12" Powerbook (as Client), open the Terminal and "talk" to the Controller (norbury15). First, we'll query to see what grids are available:

```
tiger12:~norbury15$ xgrid -h norbury15.local -grid list
```

Our syntax includes the xgrid command, -h for hostname of the controller ("norbury15.local", since I'm using an Airport wireless network here at home), and the options -grid list to give us the list of available grids on our local network. The result is:

```
{gridList = (0); }
```

If we had two grids available, we'd see this result:

```
{gridList = (0,1); }
```

But we only have one grid and from the GUI above, we know it's called the default Xgrid. Let's take a look at information we can get from the CLI:

```
tiger12:~norbury15$ xgrid -h norbury15.local -grid attributes
-grid 0
```

We see this for our result:

```
{gridAttributes = {gridMegahertz = 0; isDefault = YES; name =
Xgrid; }; }
```

The returned attributes include the current workload (gridMegahertz): 0 because there are no jobs running; whether the grid is the default one for the controller (isDefault): YES; and the name of the grid (name): Xgrid.

Let's run a job!

```
tiger12:~norbury15$ xgrid -h norbury15.local -job run /bin/sh
-c cal 2005
```

This command initiates a job that calls the default month's calendar. When we hit enter and take then run over to take a quick look at the Xgrid Admin GUI app on our controller (norbury15), we can see some information while the job is running. Here's the Overview tab with a job running:




**Figure 7: Xgrid Overview During Job Run**

Click the Agents tab:

# Innovations by InterSystems



*Rapid development with robust objects*



*Lightning speed with a multidimensional engine*



*Easy database administration*



*Massive scalability on minimal hardware*

## The Object Database Created With Several Objects In Mind.

Caché is the first multidimensional database for transaction processing and real-time analytics. Its post-relational technology combines robust objects and robust SQL, thus eliminating object-relational mapping. It delivers massive scalability on minimal hardware, requires little administration, and incorporates a rapid application development environment.

These innovations mean faster time-to-market, lower cost of operations, and higher application performance. We back these claims with this money-back guarantee: *Buy Caché for new application development, and for up to one year you can return the license for a full refund if you are unhappy for any reason.\** Caché is available for Unix, Linux, Windows, Mac OS X, and OpenVMS – and it's deployed on more than 100,000 systems ranging from two to over 50,000 users. We are InterSystems, a global software company with a track record of innovation for more than 25 years.



Try an innovative database for free: Download a fully functional, non-expiring copy of Caché, or request it on CD, at [www.InterSystems.com/Cache2HH](http://www.InterSystems.com/Cache2HH)

\* Read about our money-back guarantee at the web page shown above.

© 2006 InterSystems Corporation. All rights reserved. InterSystems Caché is a registered trademark of InterSystems Corporation. 4-06 CacheInno2MaTe




Figure 8: Xgrid Admin Agents During Job Run

And now click on the Jobs tab:




Figure 9: Xgrid Jobs Tab During Jobs Run

Finally, here is a screenshot of all the commands we've used to query the grid and initiate the job, along with the results from the job request:




Figure 10: Xgrid CLI Client on tiger12 Client Machine

Take a look at the xgrid man pages for more options and some very good examples (rare for man pages but then Apple wrote this one!).

Now...

Apple's Server Solutions site has a large section devoted to their proprietary tools for cluster computing (<http://www.apple.com/xserve/cluster/>). They provide ample information on their cluster technology solutions: G5, Xgrid, Xcode, Shark (performance optimization tool) and the Accelerate framework, as well as their cluster products: Apple Workgroup Cluster, Xserve G5 Cluster



THE LAW OFFICE OF  
BRADLEY M. SNIDERMAN

Helping clients with their software legal issues.

- Trademark and Copyright Registration
- Trade Secret Protection
- Licensing and Non Disclosure Agreements
- Assist with Software Audits

I am an attorney practicing in Intellectual Property, Business Entity Formations, Corporate, Commercial and E-commerce Law.

Please give me a call or an e-mail. Reasonable fees.

23679 Calabasas Rd. #558 • Calabasas, CA 91302  
PHONE 818-706-0631 FAX 818-706-0651 EMAIL brad@sniderman.com

Subject: HELP!

## IT DEPARTMENT

From: Joe Arnold (Graphics Department)  
Sent: Today at 10:45 AM  
To: IT Department  
Subject: HELP!

We need you to find a solution that is going to:

- Let our Graphics and Prepress people find their files quickly
- Simplify our archiving process
- and INTEGRATE with our FileMaker Pro Database

... by tomorrow.

JOE -  
check out ...  
[www.meta-comm.com](http://www.meta-comm.com)



New Version

## DigitalStorageManager<sup>TM</sup> 1.5

Intranet Search and Archive Management for Creative Workgroups

Digital Storage Manager is the first intranet search and archive management system for creative and prepress workgroups. Digital Storage Manager offers:

- Ability to search and instantly find files by production related metadata.
- A simple one-step Archiving solution including integration with: CA Software ARCserve, Rimage, EMC Dantz Retrospect, Mac OS X File Burner, Roxio Toast, Microsoft Removable Storage Manager.
- Secure Web Portal provides your customers, freelancers or other departments the ability to search and download files via the web.
- Integrates with existing job or production management systems, including FileMaker Pro.



**Download a 21 Day Trial or Call to Schedule a Demo**

Web: [www.meta-comm.com](http://www.meta-comm.com) Tel: 1-800-771-6382 Email: [sales@meta-comm.com](mailto:sales@meta-comm.com)

Node, Xserve Raid, OS X Server and Xsan. Their growing suite of integrated products provide robust and inexpensive solutions for customers who require solutions to their compute intensive application needs.

Steve Jobs' announcement in his WWDC 2005 keynote that Apple will deliver Macs with Intel microprocessors by Summer of 2006 ignited a considerable amount of discussion and speculation. The move from the PowerPC chip is controversial, no less for the apparent alliance with a company that has traditionally been aligned with "The Dark Side". This has, of course, sparked quite a bit of online discourse about running Windows on Intel-based Macs.

Apple is expected to use the Pentium M chip inside their next generation of desktop and portables because of its low power consumption and high performance. The Pentium M is geared to deliver more performance per clock cycle in order to consume less energy, which will ultimately make it run cooler than Pentium 4 chips. Of note, virtualization technology is built into Intel chips which will allow the machines to be partitioned to run different types of software like Windows or Linux at once, on top of Mac OS X. And hardware virtualization enables a system to run at near full-speed.

Apple has issued several statements saying they will not prevent Windows from running on Intel-based Macs but they will not allow the Mac OS to run on

non-Apple machines. In fact, the Macs shipping as part of the Developers Transition Kit sport a security chip called the Trusted Platform Module that contains an encrypted serial number that verifies the OS is running on Apple hardware.

An interesting aspect of this is the potential development of cluster and grid virtualization. Traditional virtualization systems like VMWare emulate a PC down to the hardware devices. In other words, VMWare simultaneously creates multiple x86 virtual computers (guest virtual machines) and each virtual machine (VM) has it's own virtual CPU, disk, memory, etc. and all the virtual hardware is mapped to your computer's real hardware (the host machine). One limitation is significant emulation overhead and little to no opportunity for optimization. The VM (virtual machine) runs multiple kernels and full installations. Running virtualization on a cluster creates an environment with automatic load balancing via process migration: the master node provides system services while the compute nodes run the application. This scales up performance. Remember the difference between clusters and grids: with a cluster, there is a single point of administration while in a grid system, there are often different operating systems in different domains. While the cluster model provides a simpler, more well defined scenario, virtualization on grids will see advances in the future and Intel-based Macs could very well fill the growing need for low cost, commodity based distributed computing solutions.

**Gas is Getting Expensive.  
Time is Gas.  
The Internet is NOT a cloud.**



**It's your Help Desk.  
(might as well be global)**

**Mac HelpMate Professional  
Zero-Configuration/Branded  
Remote Support Service.  
(a real good thing)**

**www.MacHelpMate.com**



## **About The Author**

*Mary Norbury-Glaser is IT Director at the Barbara Davis Center for Childhood Diabetes, an affiliate center at the University of Colorado Health Sciences Center in Denver, Colorado. She has too-many-years-to-count experience in cross-platform systems implementation and administration in the education sector. You can reach her at [norburym@mac.com](mailto:norburym@mac.com).*

**MACTECH**  
M a g a z i n e

**[store.mactech.com/riskfree](http://store.mactech.com/riskfree)**

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# PROVIDING PROGRESS FEEDBACK DURING SCRIPT EXECUTION

**M**any AppleScripts do not provide progress updates to the user during processing. Most of the time, when a script is run, it simply performs the appropriate tasks “behind the scenes,” so to speak. If run as an application, a script may appear in the Dock when launched. However, this hardly provides detailed information to the user about what is actually occurring. Sometimes, a script may not need to provide progress updates to the user. However, there are situations when providing such feedback is a good idea.

In this month’s column, we will walk through the process of creating a script that will provide visual progress information to the user during processing. The script we will create will save selected email messages in Mail as text files into a user-specified output folder. Since the script will have the ability to process multiple selected email messages, we will write our code to provide a visual indication to the user of which message is currently being processed. Once you learn how to provide this type of feedback, then you can begin integrating this same technique into your other scripts, making them more user friendly.

If you followed along with some of my past *AppleScript Essentials* columns, then you are probably familiar with AppleScript Studio, a feature set of Xcode and Interface Builder, the Mac OS X developer tools that come with OS X. AppleScript Studio provides a way for developers to build AppleScript-based applications, complete with robust interfaces. In this month’s column, we will use AppleScript Studio to add a progress interface to our script, complete with text feedback and a progress bar.

## Displaying a Basic Progress Interface with AppleScript

Before we get started with AppleScript Studio, let’s discuss how to provide progress information to the user in a non-AppleScript Studio-based script.

In some cases, taking the time to construct an AppleScript Studio application may not be the best solution. For example, your script may be very simple, or it may be an existing script that is too complex to warrant conversion to AppleScript Studio at this time. You may just want a quick and easy way to provide feedback to the user. In these types of situations, the easiest method is to make use of the `display dialog` command, which can be found in the *User Interaction* suite of the Standard Additions scripting addition that is installed with Mac OS X.

Using the `display dialog` command, you can configure a script to display text messages in a no-frills dialog window at various times during script execution. The following example code demonstrates how a `display dialog` command can be used to provide such feedback.

```
set theOutputFolder to (choose folder with prompt "Select an
output folder:") as string
tell application "Mail"
    set theSelectedMessages to selection
    set theMessageCount to count theSelectedMessages
    repeat with a from 1 to theMessageCount
        display dialog "Processing message " & a & " of " &
theMessageCount giving up after 1 with icon note
        set theMessageContent to content of item a of
theSelectedMessages
        set theArchivePath to theOutputFolder & "Archived
Message " & a & ".txt" as string
        set theArchiveFile to open for access theArchivePath
with write permission
        set eof of theArchiveFile to 0
        write theMessageContent to theArchiveFile
        close access theArchiveFile
    end repeat
end tell
```

### Register

**Get your .COM  
or any other  
domain name  
here!**

**FREE** with every domain:

- **FREE!** Starter Web Page
- **FREE!** Getting Started Guide
- **FREE!** Complete Email
- **FREE!** Change of Registration
- **FREE!** Parked Page w/ Domain
- **FREE!** Domain Name Locking
- **FREE!** Status Alert
- **FREE!** Total DNS Control

Just visit

**[www.mactechdomains.com](http://www.mactechdomains.com)**  
to register for your domain today!

**starting  
at  
\$1.99**

**when a non-domain name product  
is purchased. Limitations apply.**

In the previous code, the first line of the script will prompt the user to select an output folder. Next, the script will retrieve a list of selected email messages in Mail. The script will then proceed to loop through each of these messages. During each repeat loop cycle, the script will retrieve the content of the current message, and save it into a file in the output folder chosen by the user.

As you can see, in order to provide visual feedback to the user during processing, I have made use of a `display dialog` command within the script's repeat loop. This will cause the script to display a message to the user each time the script begins processing one of the selected email messages. I have configured the `display dialog` command to indicate the current message count, as well as the total count of selected messages. See figure 1.




**Figure 1. A Basic Script Progress Dialog**

I have also configured the `display dialog` command to automatically dismiss the dialog after it has been displayed for 1 second. This will ensure that the script can proceed automatically with further execution, without the user being required to actually click a button in order to proceed.

Upon the successful execution of the previous code, the output folder specified by the user should contain text files containing the contents of the selected email messages. See figure 2.




**Figure 2. Archived Email Messages**

The `display dialog` command can also be used independently, rather than within a repeat loop to indicate various tasks that are being performed by the script.

Now that we have discussed providing basic progress information to the user, let's move on to AppleScript Studio. In the remainder of this month's column, we will walk through the

process of creating a script that will perform the exact same function as the previous code, only with a more robust interface.

## Building the AppleScript Studio Project

You may recall that the first step in building an AppleScript Studio project is to create a new project in Xcode. Begin by launching Xcode.

Please note that the AppleScript Studio project covered in this article was developed using Mac OS X 10.4.2 and Xcode 2.1. Please be aware that new software versions often result in changes in AppleScript terminology. Therefore, if you are using software versions other than those that I have specified, your required terminology may differ slightly from that which I am using in this article.

Create a new project by selecting *New Project...* from the *File* menu in Xcode. When prompted, select a project type of *AppleScript Application* from the list of available project templates, enter a project name of *Archive Selected Messages*, and specify an output folder for the project. Xcode will duplicate the *AppleScript Application* project template into the specified output folder, and open it for you. See figure 3.




Figure 3. Archive Selected Messages Project Window

## Building the Interface

Next, we will create an interface for providing progress information during processing. Double click on the *MainMenu.nib* file within your Xcode project window. This will open the project's default view in the Interface Builder application.

## Designing the Window

By default, the project's interface should already contain an empty window view. This window will be the basis for our interface. Click on the window, and give the window a name by entering *Archive Selected Messages Progress* into the *Window Title* field in the *Inspector* palette. See figure 4. If the *Inspector* palette is not visible, select *Show Inspector* from the *Tools* menu in Interface Builder.




Figure 4. Preparing the Progress Window

Once you have specified the title for the window, de-select the *Visible at launch time* checkbox in the *Inspector* palette. This will cause the window to be hidden when the project is first launched.

You may want to make some other adjustments to the window configuration at this time, as well. For example, you may want to disable the ability for the user to close or zoom the window. Refer to figure 4 for the settings that I have specified for my progress window.

Next, click on the *Cocoa Controls and Indicators* tab in the toolbar of the *Palettes* window. If the *Palettes* window is not visible, select *Palettes > Show Palettes* from the *Tools* menu. Next, locate an *NSProgressIndicator* progress bar interface element in the *Palettes* window, and drag it into your project's window. See figure 5.




Figure 5. Adding an NSProgressIndicator

Next, click on the *Cocoa Text Controls* button in the toolbar of the *Palette* window, and drag an *NSTextField* into your interface window. Enter some default text, such as *Waiting to process...* into the text field's contents.

Continue to arrange and design the interface, making sure to adhere to Apple's standards for human user interface

**Really fast compilers have arrived ...**



**Make your core duet really sing!**

**New compilers, debuggers, libraries, and powerful tools  
from Intel® and Absoft are now available  
for Apple® computers using Intel® processors.**

Contact [sales@absoft.com](mailto:sales@absoft.com) or call 248-853-0050 to order today!  
<http://www.absoft.com>

**absoft®**

Development tools for Apple® customers since 1984

guidelines, which can be found in the ADC Reference Library, both online and in Xcode's documentation. See figure 6 for an example of my completed interface design.




Figure 6. Example Progress Interface Design

## Preparing the Interface for AppleScript Interaction

Once you have finished designing your progress interface, the elements that make up the interface must be prepared to interact with the AppleScript code within your project. To do this, you must assign AppleScript names to various interface elements, as well as configure certain elements of the interface to respond to event handlers.

First, we will assign an AppleScript name to the main window itself. To do this, click on the window to select it. Next, choose *AppleScript* from the popup button at the top of the *Inspector* palette, and enter the name *Progress*

into the *Name* field. See figure 7. The process of assigning AppleScript names to the other interface elements will be the same.




Figure 7. Assigning an AppleScript Name to the Progress Window

Select the progress indicator bar and text field in the window, and assign AppleScript names of *Progress Bar* and *Progress Text*, respectively.

As the last step in configuring the progress interface, we must configure the solution to run our AppleScript code when launched. This will be done by enabling an event handler. To do this, first select *File's Owner* in the *MainMenu.nib* window. A list of available event handlers will be visible on the *Inspector* palette, beneath the *Name* field. Enable the *launched* event handler by selecting its checkbox in the list of event handlers. Next, link the event handler to the AppleScript code in your project by selecting the checkbox next to *Archive Selected Messages.applescript* in the *Script* area at the bottom of the *Inspector* palette, beneath the event handler list. See figure 8.



**MacResource**  
the.mac.resource

**Wireless Products**

- Airport Cards
- Extreme: \$69.99
- Standard: \$139.99
- Routers: \$59.99

**Accessories**

- Apple Mice: \$19/\$69
- Apple Kbds: \$19/\$49
- Kensington 4-btn Mice: \$19

**Parts for Every Mac System!**

**Logic Boards**

- G3 B&W/G4 PCI : \$129
- G4 AGP Graphics: \$179
- G4 Gigabit: \$199
- G4 Digital Audio: \$239
- G4 Quicksilver: \$449
- G4 Mirrored Drive Doors: \$489
- G3 iMac Logics From: \$69
- G4 iMac Logics From: \$199
- G5 Towers and Xserves \$149-\$799

**Power Supplies**

- G4 iMac 15/17/20: \$79/\$99/\$119
- G5 iMac 17/20: \$149/\$179
- G5 Tower: \$169/\$199
- G4 Quicksilver: \$179
- G4 Digital Audio: \$189
- G4 Gigabit Ethernet: \$199
- G4 AGP Series 237w: \$219
- G4 AGP Series 208w: \$179
- G4 Mirrored Drive Doors: \$199

Logic board and power supplies require exchange.

**Thousands of parts in stock for all Mac systems!**

**www.mac-resource.com**

**Mac Systems**

- G4/867Mhz: \$699
- G4/933Mhz: \$899
- G4/1.25Mhz SOLD OUT
- G5 1.6Ghz: \$899
- G5 1.8Ghz: \$1099
- G5 1.8Ghz DP \$1249
- G5 2.0Ghz DP \$1399
- eMac 800Mhz: \$449, eMac 1GHz: \$549
- iMac G4/700 15": \$449 G4/1.25Ghz 17": \$749

**XSERVES AND RAIDS**

**BEST PRICES!!!!**


- G4 XSERVES STARTING@ \$999!!!!
- G5 XSERVES STARTING @ \$1499!!!!
- 1 TB XSERVE RAIDS FROM \$3499!!!!
- 2.8 TB XSERVE RAIDS FROM \$4499!!!!
- 5.6 TB XSERVE RAIDS FROM \$6999!!!!
- RAID CARDS, FIBRE CARDS, DRIVE MODULES
- SERVICE KITS, CONTROLLER MODULES ETC!!!!

**Displays**

**REFURBISHED DISPLAYS**

- Apple 17" Studio LCD: \$349
- Apple 23" Cinema: \$899
- Apple 22" Cinema: \$799
- Apple 15-inch Studio LCD: \$179
- Apple 17" Studio CRT,ADC or VGA: \$69.99

Products may be refurbished or demo, call for more information



# PHONEPIPE<sup>TM</sup> VOIP Service

## Residential Plans

### The PhonePipe 500

**\$14.99**

**500 Minutes - US and Canada**

\*Additional Minutes 3.5 cents

### The PhonePipe 900

**\$19.99**

**900 Minutes**—US, Canada, Austria, Belgium, Chile, China, Denmark, France, Germany, Hong Kong, Ireland, Italy, Malaysia, Netherlands, New Zealand, Norway, Singapore, South Korea, Spain, Sweden, Switzerland, Taiwan, UK

\*Additional Minutes 3.5 cents

### The PhonePipe Unlimited

**\$24.99**

**Unlimited Minutes  
US and Canada**

### The PhonePipe Unlimited International \$34.99

**Unlimited Minutes**—US, Canada, Austria, Belgium, Chile, China, Denmark, France, Germany, Hong Kong, Ireland, Italy, Malaysia, Netherlands, New Zealand, Norway, Singapore, South Korea, Spain, Sweden, Switzerland, Taiwan, UK

**\*All Plans Include Caller ID, Voicemail, and Three-Way Calling**

## Business Plans

### Base Line

Starting at

**\$20.00 a Month**

**Includes:**

- VM • CallerID
- 3-Way Calling
- Call Forwarding
- Call Waiting Plus

\*Long Distance Starting at  
2 Cents Per Minute

### Enhanced Line

Starting at

**\$30.00 a Line**

**Includes:**

- Basic Line Features
- Microsoft Outlook Integration
- Find Me-Follow Me
- Simultaneous Ring
- Personal Web Portal

\*Long Distance Starting at  
2 Cents Per Minute

### Unlimited Line

**\$49.99 a Line**

**Includes:**

- Enhanced Line Features
- Unlimited Long Distance  
to US and Canada

\*Long Distance Starting at  
2 Cents Per Minute

**\*All Business Lines Include PBX Features: Call Transfer, Music on Hold (Customizable), Call Hold, 4 or 5 Digit Dialing**

**Optional Features: Auto Attendant, Call Center, Reception Console**

**www.PhonePipe.com**  
**1-877-300-3035 Ext 8200**




**Figure 8. Enabling the Launched Event Handler**

## Adding the AppleScript code

Now that you have configured your interface, it is time to begin adding the AppleScript code into your project. Return to Xcode, and double click on *Archive Selected Messages.applescript* file in your project to begin editing the AppleScript code.

### Preparing the Launched Event Handler

Since we configured our interface to respond to the launched event handler, we will need to add this handler to our project's code. If you saved your project's interface in Interface Builder, then this code may have been automatically inserted for you within your main project script in Xcode. If it does not exist, enter it as follows:

```
on launched theObject
end launched
```

Any code that is entered into this launched handler will be executed when our project application is launched. As we proceed, enter all code below within the launched handler.

### Get the Output Folder

It is now time to begin adding the processing code to our project. Begin by entering the following code, which will prompt the user to select an output folder. Take note that this code is identical to that used in our non-AppleScript Studio example.

```
set theOutputFolder to (choose folder with prompt "Select an output folder:") as string
```

### Get the Selected Messages

Next, add the following code, which will retrieve a list of any selected email messages in Mail, and will then count the detected messages.

```
tell application "Mail"
    set theSelectedMessages to selection
end tell
```

```
set theMessageCount to count theSelectedMessages
```

## Show the Progress Window

You may recall that we configured our progress window to not be visible on launch. The following code will now make this window visible to the user.

```
set visible of window "Progress Window" to true
```

## Prepare the Progress Bar

We are now ready to begin preparing the progress bar within our window. First, we will set the maximum value property of the progress bar to the number of detected email messages.

```
set maximum value of progress indicator "Progress Bar" of window "Progress Window" to theMessageCount
```

Next, in order to ensure that our progress bar will display incremental progress, we will set the indeterminate property of the progress bar to a value of false. A progress bar with an indeterminate property value of true will appear as a blue and white striped bar, as can be seen in figure 6, and this is not desirable for providing incremental progress.

```
set indeterminate of progress indicator "Progress Bar" of window "Progress Window" to false
```

## Looping Through the Selected Messages

We will now add code to loop through the selected email messages. Enter the following code into the script:

```
repeat with a from 1 to theMessageCount
end repeat
```

## Updating the Progress Text

The following code should be added immediately within the repeat statement. It will update the text field in our interface to indicate the current message being processed. As in our non-AppleScript Studio example, this text will tell the user the current message count, as well as the total message count.

```
set contents of text field "Progress Text" of window "Progress Window" to "Processing message " & a & " of " & theMessageCount
```

## Getting the Current Message's Contents

Next, add the code below, which will retrieve the content of the current email message.

```
tell application "Mail"
    set theMessageContent to content of item a of theSelectedMessages
end tell
```

## Saving the Contents to a File

Again, using the exact code from our non-AppleScript Studio example, add the following to your script. This code will write the content of the current message to a text file in the specified output folder.

```

set theArchivePath to theOutputFolder & "Archived Message " &
a & ".txt" as string
set theArchiveFile to open for access theArchivePath with
write permission
set eof of theArchiveFile to 0
write theMessageContent to theArchiveFile
close access theArchiveFile

```

## Updating the Progress Bar

Finally, for the last section of code within the repeat loop, add the following code. This code will set the content of the progress bar to the current repeat loop increment, thus increasing the progress bar's display to accurately reflect the current number of messages processed.

```

set content of progress indicator "Progress Bar" of window
"Progress Window" to a
update window "Progress Window"

```

Please note that the last line in the preceding code will update the window, ensuring that the progress bar's interface is refreshed each time its content value is changed.

## Completing the Code

To complete the handler, add a quit command at the end of the handler, just outside of the repeat statement. This will ensure that the application quits, once processing is complete.

```
quit
```

Next, wrap all of the code within the handler inside of a try statement, configured as follows:

```

try
on error theErrorMessage number theErrorNumber
if theErrorNumber = -128 then quit
error theErrorMessage number theErrorNumber
end try

```

This try statement will trap for a user cancelled error, error number -128, which would occur if the user clicks the *Cancel* button when prompted to select an output folder.

Now that your script is complete, the following example code shows how the completed launched handler should appear within your main project script.

```

on launched theObject
try
set theOutputFolder to (choose folder with prompt
"Select an output folder:") as string
tell application "Mail"
set theSelectedMessages to selection
end tell
set theMessageCount to count theSelectedMessages
set visible of window "Progress Window" to true
set maximum value of progress indicator "Progress Bar"
of window "Progress Window" to theMessageCount
set indeterminate of progress indicator "Progress Bar"
of window "Progress Window" to false
repeat with a from 1 to theMessageCount
set contents of text field "Progress Text" of window
"Progress Window" to "Processing message " & a & " of " &
theMessageCount
tell application "Mail"
set theMessageContent to content of item a of
theSelectedMessages
end tell
set theArchivePath to theOutputFolder & "Archived
Message " & a & ".txt" as string
set theArchiveFile to open for access theArchivePath
with write permission

```



# The #1 Open Source Database Management Tool in Town — Navicat for MySQL & PostgreSQL

To receive a **20% discount**  
submit this code: Mtech06 at  
<http://www.navicat.com/mactech.html>


Navicat is the industry leading GUI Client for the most popular open source databases. The Cross Platform capability and Ease of Use of Navicat give you the flexibility and convenience to manage and develop your database. Navicat allows database developers/ administrators to be more productive by providing an intuitive graphical user interface to MySQL and PostgreSQL.

With Navicat, you can use a single, Industry's best-selling tool for both development and database object management on: Windows, Mac OS X and Linux. Managing multiple databases on your local/ remote server is easier than ever!

Navicat Migration Tool lets you import MS Access, Excel and other most popular file formats to the database. Other utilitarian tools such as Data Transfer, Data Synchronization, Backup and Report will empower you to work more efficiently.

With Navicat Batch Job Scheduling feature, you can sit back and let Navicat do the rest for you!

Find out more, and join our 1,100,000+ users at: [www.navicat.com](http://www.navicat.com).



```

set eof of theArchiveFile to 0
write theMessageContent to theArchiveFile
close access theArchiveFile
set content of progress indicator "Progress Bar" of
window "Progress Window" to a
update window "Progress Window"
end repeat
quit
on error theErrorMessage number theErrorNumber
if theErrorNumber = -128 then quit
error theErrorMessage number theErrorNumber
end try
end launched

```

## Testing the Project

Now that our project is complete, it is ready for testing. To test the project, first launch Mail, and select multiple email messages. For best results, you may want to select a large number of messages, in order to ensure that the progress bar will increment properly. Next, select **Build and Run** from the **Build** menu in Xcode. If everything works as expected, your solution should launch, and you should be prompted to select an output folder. After choosing a folder, the project's interface should be displayed, and the selected messages should be processed and saved into the specified output folder.




Figure 9. The Completed Progress Interface

## Other Options

In AppleScript Studio, a progress indicator may be displayed as a bar, as we have seen. However, a progress indicator may also be displayed as a spinner. See figure 10.




Figure 10. Example of a Progress Spinner

A progress spinner may be useful in situations where your code does not warrant providing incremental feedback to the user, yet you still wish to indicate that processing is occurring.

## In Closing

The project discussed in this month's column should give you some basic ways to implement progress feedback in your scripts. Now, it is up to you to begin making your scripts more user friendly by providing such feedback to your users.

If you have difficulty getting any of the specified code to work, or if you prefer to review the actual project files, you may wish to download the example code. I have made the sample project discussed in this article available for download from my web site at the following URL:

<http://www.automatedworkflows.com/files/demos/MacTECH.01.06.Example.zip>

Until next time, keep scripting!



## Control real-world devices

from your Mac with a wide variety of standard applications and programming languages such as:

- REALbasic
- AppleScript
- Runtime Revolution
- FileMaker
- Ragtime
- Xcode
- C/C++
- Cocoa
- Java
- CodeWarrior
- Carbon
- MaxMSP
- LabView
- DirectorMX
- 4th Dimension
- and the UNIX-Shell

Simply buy one interface from the SERVICE USB series and control motors, lamps, relays. Log values like temperature, pressure or brightness. Get the status of switches, light barriers or reed-contacts.

All interfaces ships with a complete SDK, drivers, sample apps and documentation.



SERVICE USB plus • \$465.48



SERVICE USB • \$297.80

Custom made models, USB-POS accessories and special electronics are also available.

Well suited pre- and after sales eMail-support is available as a matter of course.

[www.bkohg.com/products.html](http://www.bkohg.com/products.html)

## About The Author



Ben Waldie is the author of the best selling books "AppleScripting the Finder" and the "Mac OS X Technology Guide to Automator", available from <http://www.spiderworks.com>. Ben is also president of Automated Workflows, LLC, a company specializing in AppleScript and workflow automation consulting. For years, Ben has developed professional AppleScript-based solutions for businesses including Adobe, Apple, NASA, PC World, and TV Guide. For more information about Ben, please visit <http://www.automatedworkflows.com>, or email Ben at [applescriptguru@mac.com](mailto:applescriptguru@mac.com).

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# ADDING AJAX TO A WEBSITE

CREATING A DYNAMIC, USER-FRIENDLY WEBSITE INTERFACE IS SIMPLE  
AND STRAIGHTFORWARD

BY ANDREW TURNER

### Introduction

Modern websites and web-applications appear drastically different from sites on the web 5 and 10 years ago. Tools like GoogleMail, BaseCamp, and TiddlyWiki have revolutionized the general concept of what a webpage can do and how users interact with it. The days of clicking on a simple hyperlink to be taken to a new page, or sitting and waiting for a form submission are rapidly dwindling.

The technology driving these sites is not really new, but their application and use has only recently become widespread and supported by a majority of web browsers. Furthermore, many web developers feel daunted by the rapid pace of the changing techniques and don't have a clear understanding of how the technologies are implemented and used.

One of the most revolutionizing of these technologies has been dubbed AJAX (or Ajax depending upon whom you ask). Ajax is responsible for dynamic page content, marking database entries, and in-line text-editing without the need for page-reloads or large, complex plug-ins like Flash or Java.

The goal of this article is to teach you the basics of Ajax and demonstrate that it is not as difficult a concept as it may first appear. In reality, Ajax is simple and easy for any web developer to add to their new or already existing site.

### What is Ajax?

AJAX is an acronym for: Asynchronous Javascript and XML. The most important concept of AJAX is the "asynchronous" part. Asynchronous communication means that commands do not need to wait for a response. By contrast, synchronous communication requires the command to wait for a response before continuing. An example of synchronous communication is a typical hyperlink; the user clicks a link, and then waits while the resulting page is requested, returned, and displayed. An

asynchronous example may be having a contact name and phone number lookup with dynamic autocomplete with names already in the database. For application developers, it may be useful to think of synchronous communication as modal, while asynchronous is non-modal.

Javascript is the client-side scripting language that has been used to implement the input, output and server-response handling. Because the code is client-side it is fast and scales up with increasing usage. The last part of the AJAX acronym is XML (extensible Markup Language), which is used in the response to encapsulate information. By using a structure like XML, the client can parse the tree for specific data without having a predetermined order of the data. As we will discuss, XML or Javascript are not required to implement "Ajax" in a site.

Ajax uses several other technologies and functionality to work. XHTML and the Document Object Model (DOM) allow Javascript to dynamically modify a webpage. CSS (Cascading Style Sheets) are not necessary, but are typically employed to provide for easy layout and design of a webpage and allow the Ajax functionality to work on the data, and not the view.

The reason the technology is referred to as either AJAX or Ajax is because of the blurring between the concept, and the implementation. Ajax (non-acronym) has become the terminology associated with the ability to dynamically modify a webpage or backend content without requiring a page reload, while AJAX (acronym) is the specific implementation of Ajax employing Javascript and XML. The term Ajax was coined by Jesse Garnett of AdaptivePath (see resources) as a better name than the previously used "Asynchronous JavaScript + CSS + DOM + XMLHttpRequest". The technologies were all originally combined by Microsoft for developing their Outlook Personal Information Manager (PIM) web application interface.

## Why use Ajax?

While the web has inarguably drastically changed the way a computer user works, to date they haven't been able to fully replace, or even work entirely in tandem with, desktop applications. To clarify, a desktop application is software that must be installed on a user's computer and is run in a self-contained window/context. By contrast, a web application operates primarily within a user's browser and is not required to be installed on a machine. This provides users access to the application and associated data from any computer using a suitable browser.

However, with the advent and widespread use of technologies such as Ajax, users can now complement, or even replace, their desktop applications with a web application. Many users are now switching to reading their email in GoogleMail, storing their documents and notes in TiddlyWiki, reading their RSS news via Gregarious, or working with colleagues in BaseCamp.

Adding Ajax to your own web site or web application provides a much smoother, and rich user experience. Furthermore, Ajax websites more closely imitate their desktop counterparts, allowing users to interact and understand the user interfaces in a similar way.

Ajax is also a relatively straightforward and simple technology to provide in a website. Developers may quickly become confused by all of the terms, techniques, and options. However, at its core, Ajax is quick to setup and begin using, and completely flexible for whatever the developer and site requirements need. Ajax can be used for features such as

inline form validation, database queries, content editing, drag-and-drop, page updating, and many others

## Setting up the framework

### Parts of Asynchronous Communications

In order to understand the essential parts of an Ajax framework, we will discuss the necessary parts of asynchronous communications. The parts are split up by **Client**, the user's browser, and **Server**, the website hosting server.

**Client:** Create a request object

**Client:** Assign a response handler

**Client:** Send a query to the server

**Server:** Receive the query, and perform operations

**Server:** Send the response to the client

**Client:** Handle the server response




**Figure 1: Asynchronous communications allow a user to continually interact with the browser, and provides dynamic updating of the web site**

**Post your Project >> Receive Bids >> Choose a Developer >> Send Payment >> Done.**

# MacDeveloper

**Project Outsourcing. Simplified.**

<http://macdeveloper.net>

The client requests to the server can happen continually, updating the web page or application on each response received. Each request happens separately from the interface, allowing the user to continue to view and interact with the web page.

However, it is a good idea to only support a single asynchronous command at a time as the response may affect the interface data. If multiple asynchronous requests are supported, you must be careful to handle potential conflicts due to user interaction with outdated data.

## Create a request object

The first thing to do is to create a *constructor* that will build a client-side request object. A request object is responsible for wrapping up the actual request, response handler, and state of the request.

Remember that this Javascript code is being run on the user's desktop browser. Therefore creating the request object is the one place where browser specific code is required. In this case, Microsoft's Internet Explorer uses an ActiveX object as the request object, whereas the other browsers all support an *XMLHttpRequest()* constructor call. We can interrogate the browser to find out what type it is and create the appropriate object. This function is universal for any Ajax use.

## AjaxFramework.js

```
// request object constructor
function createRequestObject() {
    var ro;
    var browser = navigator.appName;
    if(browser == "Microsoft Internet Explorer"){
        ro = new ActiveXObject("Microsoft.XMLHTTP");
    }else{
        ro = new XMLHttpRequest();
    }
    return ro;
}
```

You should now create a global request object that will be used by the client for all future communication.

## AjaxFramework.js

```
// global request object
var http = createRequestObject();
```

## Assign a response handler and handle the response

Our second step is to assign a response handler. A handler is the function that will be called when the request comes back from the server to the client's computer. This function is responsible for verifying the state of the answer, and parsing the response as appropriate. This function is implemented on a project specific basis. It needs to know what the expected response from the server looks like and how to place that response back into the user's browser document.

## AjaxFramework.js

```
// callback function that handles any state changes of our
// request to the server
function handleResponse() {
    if(http.readyState == 1){
        // request loading
        document.getElementById("status").innerHTML
            = "requesting...";
    }
    else if(http.readyState == 4) {
        // request complete
        if(http.status == 200) {
            // OK returned
            var response = http.responseText;

            // Add more advanced parsing here if desired

            document.getElementById("responseArea").innerHTML
                = response;
        }
        else
        {
            document.getElementById("status").innerHTML
                = "error: " + http.statusText;
        }
    }
}
```

The first thing the *handleResponse* function does is check the current state of the request object. If the object is loading (1), then the user is alerted to this, or if the request is complete (4) then we handle the response. This

**HostedStore™**

- HostedStore provides a **TURNKEY SOLUTION FOR CREATING A WEB STORE**, yet is extensible so developers can use the application as a base for creating customized solutions for their clients.
- The extensibility of HostedStore provides a structure which allows developers to create modules that add new functionality or alter the built-in features.
- We are seeking software resellers, hosting providers, developers, and designers that currently offer or plan to offer ecommerce solutions to their clients. Resellers receive a discount off of the published list price of our software and are free to bundle our software with their own products and/or services.

[www.hostedstore.com](http://www.hostedstore.com)

**Lasso**

# LOOK FAMILIAR?



## Get Rid of Your Fax Machine!

Receive your faxes directly to your email account  
as PDF file attachments

Service available in over 150 area codes

To check out special offers  
for MacTech readers, visit  
[www.MaxEmail.com/MacTech](http://www.MaxEmail.com/MacTech)

**maxemail®**

example just puts the response text (use *responseXML* for an XML response from a server) into our document's *responseArea*.

## Send a query to the server

Now that we have setup the request object structure, as well as the state handling function, the next step is to create a function that our webpage will be calling for each outgoing request. This function could either accept information via an input parameter, or retrieve user input by querying the document.

Once the user input is received, we create a *GET* request to a URL. It is important to note that due to security concerns the request can only be made to a server that is hosting the webpage. The domain name must be exactly the same as the request URL, if there is a preceding *www.* to the domain name. As usual, however, there are some fairly straightforward work arounds for getting external data for your Ajax requests. Several options will be discussed.

This example demonstrates using a REST input (parameters passed via the URL), but other remote query and command options are also possible. Furthermore, the *open* command supports passing a username and password to the server for accessing protected services.

### AjaxFramework.js

```
// function for filling out and sending a request - called by the
actual webpage
function sendRequest() {
    var query =
```

```
document.getElementById("queryInput").value;
    var queryURL =
"http://localhost.com/service.php?q=" + query;
    http.open('GET', queryURL);
    http.onreadystatechange = handleResponse;
    http.send(null);
    return true;
}
```

We have now completed the necessary parts of our Javascript code to handle creating, sending, and receiving an asynchronous request through a client's browser.

## Server handling of the request

The client makes a request to some service or page that is served on the same domain as the original webpage. This service for this example is expecting a value passed via the URL in the GET parameters. The response can be well formed XML, or simple text that will be parsed by the client's browser as discussed above in the *handleResponse()* function.

### service.php

```
<?php
$query = $_GET['q'];
$response = some_service_handling($query);
echo $response;
?>
```

This server page just passes the query onto another php function and then echoes the response. Since our Ajax request from the browser has made a *GET* request, this operates like any normal opening a page in a browser. However, instead of the page showing up in a window, it is handled by the client's *handleResponse()* function.

## Using a remote service

As we mentioned earlier, security does not allow the Ajax, specifically XMLHttpRequest, to call another domain in the GET URL. The way around this is provide a locally served wrapper to the remote service. We can parse and pass on each of the incoming parameters. Also, many hosting services don't allow a URL to be opened via the *fopen()* command, so this example uses *curl* to make a request to a server. The subsequent response is read by the local server and then returned to the calling Ajax function.

### remote\_service.php

```
<?php
$remote_params = "";
foreach($_GET as $key=>$value)
{
    $$key = $value;
    if($value != '')
        $remote_params .= "&".$key."=".$value;
}
$remote_url =
"http://remotehost.com/remoteservice.php?";

function get_content($url)
{
    $ch = curl_init();
    curl_setopt ($ch, CURLOPT_URL, $url);
```



```

curl_setopt ($ch, CURLOPT_HEADER, 0);

ob_start();

curl_exec ($ch);
curl_close ($ch);
$string = ob_get_contents();
ob_end_clean();

return $string;
}

$content = get_content
($remote_url,$remote_params);
echo $content;
?>

```

This example makes no checks on the incoming request. The query and parameters are passed directly onto the remote service. In a real application, it would be responsible to do some basic parameter checking before passing on the request to someone else's hosting service.

That said, it is still a means by which to provide asynchronous services in your own website. You should also be aware that making these remote calls may have longer response times. While this situation is an excellent reason why an asynchronous interface provides a better user experience, users may be left wondering if their request was just lost. Therefore, you should, when appropriate, let the user know that the request is pending in some way.

Furthermore, it would be possible to setup a timeout timer for each request that would call *abort()* on the request object if the request took too long.

## Supporting non-Javascript functionality

This framework will generally work for any modern, Javascript capable browser. However, not all users are using Javascript capable browsers, and other users may have disabled Javascript. Therefore, it is advised that your site support a non-Javascript version of your interface. At the very least, alert the user that they will not be able to use all of the functionality of your web application or page.

To provide a non-Javascript interface only when necessary, your page should use the *<noscript>* tags paired with any *<script>* sections.

## Using the Framework

The Javascript framework is logical backend functionality of an Ajax enabled website. In order to use Ajax, the page must be properly constructed and typically web developers also want the page to look nice. For both of these requirements, we will use XHTML and CSS respectively.

## Example query and response

Lets illustrate the Ajax framework with an example. Our service could return a name and phone number of a contact from our webserver. The query parameter, *q*, could be some search term, and the response would be the contact's name, and phone number. Testing such a service is easy:



# Looking for a great selection of cables?

MacTech has selected a quality vendor, with quality products for your cable needs. Benefits include:

- Lifetime warranty on cables
- Same day shipping if you order by 4 PM ET
- Express shipping orders can be ordered up to 6:30 PM ET for same day shipping
- Huge selection - 3500+ items and growing

Check out the full selection at the MacTech Cable Store

<http://www.mactech.com/cables>




`http://localhost.com/service.php?q=Jones`

We will then expect a response like:

Edward Jones, 800-555-1212

Our *handleResponse* functionality need to be expanded to *split* the response with the comma (or multiple commas for more data).

**AjaxFramework.js** (handleResponse)

```
var data = response.split(',');
```

// more advanced parsing

```
document.getElementById("contact_name").innerHTML  
= data[0];
```

```
document.getElementById("contact_number").innerHTML  
= data[1];
```

This example response illustrates how easy Ajax is to begin to use. There is no need for complex XML parsing and handling. Any simple response can be used to dynamically update web content. We must be careful, however, as we have forced the response to return the information in a specific order and format.

A more robust application should use XML to provide multiple contact data. The response handler could then iterate through the elements of the contact entry without having to predetermine the order of the response. In our example, if we switched the contact name and contact number order the application would behave incorrectly.

In this case, we would instead get the *responseXML* and parse the XML document tree similar to the DOM of the browser document.

**AjaxFramework.js**

```
var response = http.responseXML;  
var contact_name =  
    response.getElementsByTagName('name').item(0);  
var contact_number =  
    response.getElementsByTagName('number').item(0);
```

However, XML is not necessary, and may be daunting when first starting to use Ajax, or integrate it into already existing web services. Therefore, we will continue our example using the simple text response. Since the XML handling is encapsulated in the *handleResponse()* function, it is possible to later change to using XML without modifying the rest of our framework.

## Example page

To use the Ajax searching, we will need to provide an XHTML user interface for the query input, and the service response. The first thing we need to do is include our Javascript framework code in our page:

```
<html>  
<head>  
<script type="text/javascript"  
src="AjaxFramework.js" charset="utf-8"></script>  
</head>
```

Next we create the query input and "Send" button. Note the use of the ambiguous anchor link, #, in the *href* tag. We use an *href* link to allow standard style formatting of the "Send" button to match the rest of the sites hyperlinks. By using the local anchor, but with no actual anchor, the hyperlink won't cause a page refresh since the browser thinks it is just scrolling down the current page. Another option would have been to use a generic div and provide a unique formatting for Ajax link as compared to actual hyperlinks.

```
<body>  
<input type="text" size="30" id="queryInput"  
value="" />  
<a href="#" onClick="sendRequest();">Send</a>  
  
<div id="status">&nbsp;</div><br/>  
  
<div id="contact_name">&nbsp;</div>  
<div id="contact_number">&nbsp;</div>  
</body>  
</html>
```

When the "Send" link is pressed, the *queryInput* text input is sent as a query to our name lookup service. The user is free to continue to use the web browser. When the response is sent from the server, the retrieved name and number are placed in the *contact\_name* and *contact\_number* divs.

A more advanced version of this application could add in-line searching of the contact name as the user types, similar to autocomplete.

## Summary

Ajax is quickly transforming websites from repositories of data into dynamic and useful web applications. This article demonstrated how easy it is to get started with Ajax and add it to your own site. Some examples you can use it for include form checking while the user is entering information, site/document search, database row updating, or editing web content in place.

For more advanced applications you may want to look at several available and supported Ajax toolsets that provide a ready framework and lots of other functionality. Prototype (see resources) is used in Ruby on Rails for its Javascript Ajax functionality, and Sajax is an Ajax toolset for PHP code.

## Resources

### Ajax technology

Adaptive Path:

<http://www.adaptivepath.com/publications/essays/archives/000385.php>

# The iPod Showcase/Marketplace

## Protection Pack

Now Available for MacBook™ Pro

Includes:  
Keyboard cover  
Hand rest



Keyboard cover  
• Protects screen  
• Doubles as  
screen cleaner

Hand rest  
• Protects your  
notebook while typing  
• Easy to remove and reapply

You didn't compromise on your notebook—  
don't compromise on its protection.

[www.MARWARE.com](http://www.MARWARE.com)

## Be the one with all the answers.



MacTech is already read every month by tens of thousands of readers. MacTech readers represent the very heart and soul of the the Mac community. Join the crowd and sign up for MacTech today!

[store.mactech.com/riskfree](http://store.mactech.com/riskfree)

## What's the *difference*?




Brand new. It's perfect.  
It's beautiful...

and it's *vulnerable*!

It's naked, exposed to  
the elements.

It's...



and it *will*  
get scratched.



Brand new. It's perfect.  
It's beautiful...

and it's **protected**!

Invisible, durable, &  
rugged protection.

**SAVE**  
the iPod

[www.theinvisibleSHIELD.com](http://www.theinvisibleSHIELD.com)

Mozilla Ajax documentation:

<http://developer.mozilla.org/en/docs/AJAX>

XMLHttpRequest documentation

<http://documentation:>

[developer.apple.com/internet/webcontent/xmlhttpreq.html](http://developer.apple.com/internet/webcontent/xmlhttpreq.html)

## Ajax applications

TiddlyWiki: <http://www.tiddlywiki.com>

Gregarius RSS Reader: <http://gregarius.net>

Basecamp: <http://www.basecamp.com>

Geocode lookup (with source):

<http://highearthorbit.com/projects/geocode/>

## Ajax toolsets

Prototype, a Javascript framework for web apps:

<http://prototype.conio.net/>

Sajax, Ajax for PHP: <http://www.modernmethod.com/sajax/>

## Ajax Framework

The following files are the summation of the framework code developed in the article above. It can serve as a skeleton for building your own Ajax applications. Place these files in your /Library/WebServer/Documents directory on your Mac, and turn on "Personal Web Sharing" in the "Sharing Preference Pane".

### AjaxFramework.js

```
function createRequestObject() {
    var ro;
    var browser = navigator.appName;
    if(browser == "Microsoft Internet Explorer"){
        ro = new ActiveXObject("Microsoft.XMLHTTP");
    }else{
        ro = new XMLHttpRequest();
    }
    return ro;
}
var http = createRequestObject();

function handleResponse() {
    if(http.readyState == 1){

// request loading
        document.getElementById("status").innerHTML
            = "requesting...";
    }
    else if(http.readyState == 4) {

// request complete

        if(http.status == 200) {

// OK returned
            var response = http.responseText;

            document.getElementById("status").innerHTML
                = "loaded";

            document.getElementById("responseArea").innerHTML
                = response;
        }
        else
        {
            document.getElementById("status").innerHTML
```

```
                = "error: " + http.statusText;
            }
        }
    }

function sendRequest() {
    var query =
document.getElementById("queryInput").value;
    var queryURL = "service.php?q=" + query;
    http.open('get', queryURL);
    http.onreadystatechange = handleResponse;
    http.send(null);
    return true;
}
```

### AjaxDemo.html

```
<html>
<head>
<script type="text/javascript"
src="AjaxFramework.js"></script>
</head>
<body>
<noscript>
Your browser does not support Javascript. Please
upgrade your browser or enable Javascript to use
this site.
</noscript>

<input type="text" size="30" id="queryInput"
value="" />
<a href="#" onClick="sendRequest();">Send</a>

<div id="status">&nbsp;&nbsp;&</div><br/>

<textarea rows="20" cols="70" id="responseArea"
value="" ></textarea>

</body>
</html>
```

```
service.php
<?php
echo $_GET["q"];
?>
```

**MT**

## About The Author

*Andrew Turner is a Systems Development Engineer with Realtime Technologies, Inc. ([www.simcreator.com](http://www.simcreator.com)) and has built robotic airships, automated his house, designed spacecraft, and in general looks for any excuse to hack together cool technology. You can read more about his projects at [www.highearthorbit.com](http://www.highearthorbit.com).*

**MACTECH®**

**Get MacTech delivered to  
your door at a price FAR  
BELOW the newstand price.  
And, it's RISK FREE!**

**[store.mactech.com/riskfree](http://store.mactech.com/riskfree)**

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

## CRYPTOCARD'S CRYPTO-SERVER 6.3 FOR OS X:

### ELIMINATE INSECURE STATIC PASSWORDS

It is no secret that static passwords are the weakest link in the security chain, but until recently, there really was no Mac-specific alternative. That all changed when authentication technology vendor CRYPTOCARD released its first two-factor authentication solution for OS X Panther at MacWorld 2004 – winning a MacWorld “Best of Show” award in the process.

CRYPTOCARD has now launched a new version of its CRYPTO-Server authentication solution for OS X Tiger. Again, the basic premise of this technology is simple – it replaces inherently weak static passwords with secure two-factor authentication. To log on to a protected network or resource, a user must combine their security PIN (something only they know) with a one-time passcode that is randomly generated by their token for each logon (something only they have).

The new version of CRYPTO-Server does a good job of leveraging Tiger's robust support for smart card environments, but users can also opt for PIN pad tokens, key chain tokens, or software tokens. Each form factor offers unique advantages and characteristics, enabling organizations to tailor their authentication solution according to their own needs. Hardware tokens feature field-replaceable batteries that can be swapped in-service to extend device lifespan indefinitely.

We are pleased to report that our experience testing the new Tiger product was a good one. The CRYPTO-Server package contained everything required to set up the solution, and the instruction manual was clear, accurate, and easy to follow. The product emphasizes ease-of-use and tight integration with Apple's Open Directory LDAP services and as a result, installation was straightforward and relatively painless. There are also features that will simplify implementation in a real-world environment, such as a self-enrolment component called CRYPTO-Deploy, which

enables users to remotely assign and activate their hardware tokens via a Web page.

Once the CRYPTO-Server install is completed, a user will need to install the CRYPTO-Console module, an intuitive Graphical User Interface (GUI) which provides the management interface to CRYPTO-Server. CRYPTO-Console enables administrators to manage tokens, users (in non-LDAP deployments), and groups, while also providing server licensing, system configuration, and reporting functions.

The CRYPTO-Console interface is well thought out and easy to navigate, providing administrators with screens for viewing/editing users, tokens, containers, objects, and attributes. Search functions accept regular expressions for ease of use and the GUI architecture is logical and intuitive. Detailed management options are available by highlighting an object, and then Ctrl-clicking it to display drop-down menu items.

The solution appears to be extremely flexible, and can easily be enhanced and expanded with a variety of agents and plug-ins that extend strong two-factor authentication to existing Web, mail, and other security appliance infrastructure nodes. For

example, the CRYPTO-Logon for Mac OS X component makes it easy for Mac users attempting to gain secure LAN, Web, or remote access to authenticate themselves by simply inserting their smart card and entering their PIN.

All CRYPTO-Server tokens generate a unique password for every logon attempt, which makes stolen credentials useless to hackers, while simultaneously ensuring Tiger and Panther users do not have to memorize complicated credentials. CRYPTOCARD points out that this can significantly reduce the help-desk costs associated with password management while simultaneously eliminating the obvious security risks of “shoulder surfing” and users writing down their passwords.



CRYPTO-Server is also the first solution we have tested that supports two-factor authentication for Apache Web servers via its CRYPTO-Web component. (If you know of others, please let us know!) Using CRYPTO-Web we were able to secure a website, and then authenticate to it with a configured token. CRYPTO-Web should make it a simple process for administrators to secure websites by requiring users to authenticate with their token in order to gain access. Companies can also leverage out-of-the-box interoperability with network entities that provide native RADIUS support.

Unlike CRYPTOCARD's original OS X offering, which only provided client side authentication, the latest version of CRYPTO-Server for OS X also provides enterprise-ready functionality like "High Availability" which utilizes real-time multi-master replication functionality to ensure there is no single point of breakdown by switching to a replica server in the event



of system failure. This is important as it means that the authentication solution can now meet the security needs of any sized organization.

Another unique feature of CRYPTO-Server is that it offers cross-platform capability. This is important news for the majority of organizations that employ heterogeneous network environments in which any combination of Windows, Linux, or OS X servers can support any client/end-user systems running on any of the three platforms.

Other useful CRYPTO-Server features include RSA migration functionality that enables RSA SecurID DES tokens to be imported into the CRYPTO-Server, and CRYPTO-Kit, a software developer's kit that provides developers with the tools required to integrate CRYPTOCARD's technology with existing security applications/systems.

We found CRYPTO-Server for OS X to be very well thought out. Documentation is simple to follow, and the product does a good job of supporting authentication requirements, including a full compliment of token form

factors which should make it simple for any sized organization to customize an authentication solution to meet security requirements. The technology makes system configuration simple for administrators, while the familiar ATM-style logon process is easy for users to grasp.

CRYPTOCARD was the first authentication vendor to provide real two-factor authentication for the Mac, and we found that the latest version of its technology gives the company a good basis to claim leadership in the OS X authentication marketplace.



CRYPTO-Server for OS X is available in a "Five-User Kit," which includes full server software, five tokens of the user's choice, and 30 days support, for \$499. This compares favorably with other similar products from other large, well-established vendors. The innovative all-you-need-in-one-box format also makes it simple for an organization to build their security solution as required. CRYPTOCARD offers a free trial download of the CRYPTO-Server technology on its website at [www.cryptocard.com](http://www.cryptocard.com).

CRYPTOCARD Corp.  
340 March road  
Suite 600  
Kanata, Ontario. K2K 2E4  
Canada  
Phone: : North America 800-307-7042,  
International +1-613-599-2441  
Fax: +1-613-599-2442  
Web: [www.cryptocard.com](http://www.cryptocard.com)  
E-mail: [info@cryptocard.com](mailto:info@cryptocard.com)





# Other World Computing

Serving the Mac Universe since 1988

[www.macsales.com](http://www.macsales.com) • 800.275.4576

## Highly Reliable, High-Performance, Plug & Play - FireWire & USB2 External Storage Solutions from \$99.95

Add up to 1.5 Terabytes - 1500GB - for you Data, Graphics, Audio/Visual, Music, and Storage Needs with a top rated Mercury™, Neptune™, or NewerTech™ Storage Solution from OWC.

- Latest Oxford Chipsets • Top Hitachi and Seagate Drives
- EMC Retrospect Backup Certified
- 2 Year Warranty



**ATTENTION!**

**NEW! The FIRST 750GB single drive solutions!**

OWC Mercury Elite-AL Pro up to 500GB & OWC Mercury Elite-AL 800 Pro RAID 80GB to 1.5TB(1500GB) from **\$119.99**


**LAYERS**



See the full line of OWC FireWire solutions, as well as solutions by LaCie, EZ Quest, SmartDisk, and Webtech online at [macsales.com/FireWire](http://macsales.com/FireWire)

NewerTech miniStack™ V2 FireWire+USB2 Solutions with integrated FireWire and USB2 Hubs bring high performance storage and port convenience with solutions from 80GB to 750GB from **\$129.00**

**miniStack™ V2**



**NOW UP TO AMAZING 750GB Plus USB & FireWire Port Expansion The Ultimate Storage Solution**

## Mac Improvement NuPower Laptop Batteries

**newer technology**

Batteries that Run Longer and Last Longer! Built in the USA and built right for up to 56% more runtime vs. your original Apple stock battery!

PowerBook G4 Ti from **\$119.99**  
PowerBook G4 AL 12/15/17" from **\$119.99**  
iBook G3/G4 from **\$99.95**

Call or Visit [macsales.com/NewerTech](http://macsales.com/NewerTech)

## Laptop Screen Protectors

Protect your screen! There's an OWC Laptop Screen Protector (LSP) product for your Mac.

PowerBook G4 17" **\$17.99**; PowerBook G4 15" **\$17.95**  
PowerBook G3 15" **\$14.99**; iBook/PowerBook G4 12" **\$13.95**

The OWC LSPs are precision cut, glove soft leather protectors that prevent potentially permanent marks which can occur from the trackpad and keyboard while your laptop is closed.



## Network Adapters

**SONNET**

Sonnet Presto 10/100/1000 Gigabit Ethernet PCI **\$85.99**

## PRAM Batteries

Is your Mac forgetting what time it is? OWC PRAM batteries starting at **\$4.99**

## Wireless Mouse

Logitech Cordless 'Click' Optical Mouse for USB **SPECIAL \$15.99**

## The Latest Enhancements

**Rain Design i360°**

A Turntable for your iMac G5 17" and 20" **\$39.00**

**Village Tronic VTBook**

Add another CRT or Flat Panel Display to your Powerbook **\$246.99**

**iLugger iMac cases**

for the iMac G5 or for Mac mini and/or up to 20" LCD Display 5 color combinations starting at **\$99.95**



Eye candies for your iMac. Six to choose from.



# Own the Future TODAY!

[www.MacSales.com/firewire](http://www.MacSales.com/firewire)



OWC Mercury On-the-Go Portable 40GB to 160GB from **\$109.99**



## Value Done Right!

Neptune FireWire 80GB to 500GB solutions from **\$99.95**



OWC Mercury Elite Pro 'Classic' 80GB to 500GB from **\$125.99**



Backup by **EMC Retrospect**

Solutions include all cables, EMC Retrospect Backup, Intech HD Speedtools and are preformatted with free bonus software loaded all ready to plug and play your PowerMac, PowerBook, iMac, iBook, or eMac too.

## SuperDrive Your Mac From Only \$47.99!

Make Music, Movies & More! Add a Fast SuperDrive to your Mac to Burn CDs, DVDs, even 8.5GB Dual Layer DVDs.



Trade In Rebates!



Not sure what upgrade is best for your computer?

Visit: [macsales.com/MyOWC](http://macsales.com/MyOWC)

A custom shopping experience to make your Mac a better Mac.

## Speed it up as high as G4/2.0GHz or Dual 1.8GHz!

OWC Stocks the full line of G3 & G4 Processor upgrades by these leading manufacturers:



G4 Single Upgrades from \$159.00; G4/1.6GHz only \$229.00

G4 Dual/1.6GHz from \$399.99; Dual 1.8GHz from \$595.00

G4 Upgrades for PowerMac G3s, PowerMac G4s, Cube G4, PowerBook G3s - Even Legacy PowerMac 7200-9600 Models!



30 Day Money-Back on NewerTech and OWC brand upgrades!



Give us a call or check out our website. Our compatibility guide will show just what options are right to make your Power Mac, PowerBook, iMac, etc. - a Faster Mac today! Call 800.275.4576 Visit [macsales.com/faster](http://macsales.com/faster)

## Upgrade Your Memory & Save

Make your computer faster for less with Other World Computing. Memory is the most inexpensive way to maximize your Mac. From min to the Max - OWC has the right Memory for your Apple.

For NEW 2006 MacBook Pro 15.4" & 17", iMac Intel, & Mac mini Intel models:

PC5300 DDR2 667MHz 200 Pin

512MB Modules only \$49.99

1.0GB Modules only \$112.99

2.0GB Match Set only \$219.99



For PowerMac G5 and iMac G5 Models:

512MB Module from \$45.99

1.0GB Module from \$87.99

1.0GB Matched Pair from \$89.99

2.0GB Matched Pair from \$175.00

2.0GB Module from \$209.99

4.0GB Matched Pair from \$415.00

For all PowerBook G4 Models; iBook G3/G4s Models:

256MB Modules from \$26.99

512MB Modules from \$39.99

1GB (1024MB) Modules from \$109.99

Lifetime Advance Replacement Warranty.

For all PowerMac G4, eMac G4, and iMac G4:

256MB from \$24.99

512MB from \$45.99

1GB from \$87.99

## Hard Drives

Bulk up your computer by giving it higher capacity to perform for your needs.

### Hard Drive Controllers

Hard Drive Controller Cards

ACARD 2-Channel PCI SATA \$59.99

SIIG Serial ATA 4-Channel PCI-M \$79.95

PNP Tempio ATA & SATA Mac PCI Controllers

Tempio Trio FireWire/USB2/ATA-133 all in one PCI \$149.00

FirmTek 4-Channel SATA Controllers

from \$119.95

SATA Internal Mounting System

PNP Sonnet G5 Jive, add 3 SATA HDs to PowerMac G5 \$95.00

### Internal Hard Drives

For iMacs, eMacs & PowerMacs

3.5" Plug & Play 40GB to 500GB from \$47.99

250GB Super Value \$99.95

For PowerBooks, iBooks & Mac minis

2.5" 40GB to 160GB from \$69.99

100GB 7200RPM NoteBook Performance \$219.99

SATA Enclosures

OWC Mercury Elite SATA RAID Solutions

Perfect for Mass Storage or RAID. Cables included. 1 Year OWC Warranty. Starting From \$79.95

FirmTek Dual-Bay Hot Swap External Serial ATA Enclosure

with SATA PCI Controller Card \$257.95

PNP Sonnet Fusion 4-Bay SATA Kit \$549.00

Not sure what your Mac takes?

Buy with Confidence, call or use our online guides.

[macsales.com/harddrives](http://macsales.com/harddrives)

## Software

Apple OS X 'Tiger' \$99.00

full retail box version

OS X 10.2, 10.3 from \$17.99



Apple iLife '06

Make the most out of your digital life. Share the magic of your everyday with iLife '06. Only \$79.00

NEW! iLife '06 \$79.00

## iPod Replacement Batteries

newer technology

### iPod Replacement Battery Kits

Easy to Install, Tools Included + Online Installation Videos. Get up to 78% more capacity & 20+ Hours Runtime!

**NuPower**

iPod Batteries for nearly every Apple iPod

Starting From \$14.99

Not comfortable opening your iPod?

For \$39 + the cost of the battery, OWC installs it for you - iPod shipping Box and FedEx Overnight covered. to and from [macsales.com/iPodinstall](http://macsales.com/iPodinstall)

OWC's full line of iPod Batteries, Accessories, and more online at [macsales.com/iPod](http://macsales.com/iPod)



**FasterMac.net**  
Pay less. Get more. Surf faster!

Mac-Only Internet from only \$5 per month!

High-Speed Nationwide Dial-up and DSL Services  
Toll-Free Tech Support & More from Mac Experts

Visit [FasterMac.net](http://FasterMac.net) or call toll free 800-869-9152 to learn more or to sign up.

# Web Benchmarking 101

A guide to stress testing your website with ApacheBench and JMeter

By Jin Lin and Emmanuel Stein

## Introduction

Web performance testing can be a daunting subject for the uninitiated. With the sheer number of tools available and their various feature sets, it can be hard to know where to begin. In this article we will explore the use of two popular open source tools, ApacheBench (aka ab) and JMeter, to perform simple benchmarking tests. Rather than going through the extensive feature set offered by these applications, particularly JMeter, we will focus on measuring the response time and the effective throughput of your web application across different user load scenarios. Whether serving a few static pages or a fully dynamic site, the tools discussed will offer invaluable insight into the overall performance of your web applications, and serve as a basis for isolating system bottlenecks and conducting capacity planning for your web applications.

## Benchmarking Considerations

Although, the tests and tools that we will use in this article represent integral components used in the testing process, it is equally important to consider the environment in which such tests are performed. If you are running a popular website with significant traffic, you may consider performing benchmarks during non-peak hours to obtain more accurate results and avoid rendering your server unusable to users. If possible, you should test your web server in a staged environment and before actual deployment. This will enable you to adjust server parameters, install cache mechanisms, and otherwise assess your site's performance without interrupting service.

## ApacheBench

ApacheBench, a default component of the Apache Server distribution, is a simple command line tool, designed to measure the performance of Apache server. Preinstalled on Mac OS X, it is able to supply information such as elapsed time, requests per second, time per request, compression rate, transfer rate, and connection time under different concurrence scenarios.

Though not an exhaustive list (see the ab man page for more details), the following represents the most commonly

used switches employed when testing typical usage scenarios with ab.

- -n: Specifies the number of requests to perform for a given test. Use this option to measure the time required to process a given number of client requests. When used alone, without the -c option, ab does not process requests concurrently.
- -t: Specifies a timeframe, in seconds, to spend benchmarking. This option is useful in determining the number of requests that your server is able to process in a given amount of time.
- -c: Specifies number of simultaneous connections. Of use in testing how your server performs when multiple clients are hitting a given page at the same time. This switch may be used with either -n or -t, to measure the total elapsed time required to process a given number of client requests simultaneously, and the number of concurrent requests that may be processed over a set amount of time, respectively.
- -g: Exports a TSV (Tab Separated Values) file, which can be imported into applications like Excel, GNUplot, and Mathematica for further analysis and/or graphic visualization.
- -A: Used to supply username and password information for sites that require authentication. This enables you to conduct testing on password-protected sites.

Having covered some of the basic parameters of ab, we will now explore its use in the context of simple web benchmarking scenarios.

For our first example, we will be measuring the time it takes to process 1000 requests (i.e. -n 1000), from 10 simultaneous user connections (i.e. -c 10) against OS X's default user website. Make sure Apache is running by launching System Preferences and enabling personal web sharing from the **sharing preference pane**. Once Apache is running, launch the Terminal application and enter the following command: `ab -n 1000 -c 10 http://localhost/~username` to begin

the benchmark. When the command is executed the terminal will display the progress of the benchmark run, as well as, a series of metrics (Figure 1).

```

macverse:~ estein$ ab -n 1000 -c 10 http://localhost/~estein
This is ApacheBench, Version 1.3d <Revision: 1.73> $> apache-1.3
Copyright (c) 1996 Adam Twiss, Zeus Technology Ltd, http://www.zeustech.net/
Copyright (c) 1998-2002 The Apache Software Foundation, http://www.apache.org/

Benchmarking localhost (be patient)
Completed 100 requests
Completed 200 requests
Completed 300 requests
Completed 400 requests
Completed 500 requests
Completed 600 requests
Completed 700 requests
Completed 800 requests
Completed 900 requests
Finished 1000 requests
Server Software:      Apache/1.3.33
Server Hostname:      localhost
Server Port:          80

Document Path:        /~estein
Document Length:      276 bytes

Concurrency Level:    10
Time taken for tests:  0.328 seconds
Complete requests:     1000
Failed requests:        0
Broken pipe errors:     0
Non-2xx responses:     1003
Total transferred:     436305 bytes
HTML transferred:      276826 bytes
Requests per second:   3048.78 [#/sec] (mean)
Time per request:      3.28 [ms] (mean)
Time per request:      0.33 [ms] (mean, across all concurrent requests)
Transfer rate:         1330.20 [Kbytes/sec] received

Connection Times (ms)
  min  mean[+/-sd] median max
Connect:  0    0  0.1   0    4
Processing: 1    3  1.6   2   15
Waiting:  1    2  1.6   2   14
Total:    1    3  1.7   2   15

Percentage of the requests served within a certain time (ms)
 50%  2
 66%  3
 75%  3
 80%  3
 90%  3
 95%  5
 98%  11
 99%  14
100% 15 (last request)
macverse:~ estein$

```

**Figure 1. Output of `ab -n 1000 -c 10 http://localhost/~estein`**

Be aware that when you run the `ab` command against a domain or IP address the benchmark defaults to the site's home page; in this case, `index.html`. To specify a particular page, simply point to its full path (e.g. <http://localhost/~estein/MyDynamicContent.php>). This will allow you to test the performance of specific pages hosted by your site in terms of the amount of stress they each place on the server. This is particularly important when dealing with sites that have both static and dynamic content. When testing dynamic pages keep in mind that `ab` reports a failed request when it encounters content that changes over time (e.g. dynamic data). As such, when running `ab` against dynamic portions of your site, you can safely ignore the "failed requests" messages.

For our next example, we will be measuring the number of requests that can be completed within 60 seconds (`-t 60`) with 10 concurrent connections (`-c 10`). Enter the following command in the Terminal to execute this benchmark: `ab -t 60 -c 10 http://localhost/~username`. The output of this

command, seen in Figure 2, differs from the first example in that `-n` is not specified. When `-n` is not specified, `ab` sets the number of requests to a default value of 50,000. It is important to realize that, when using the `-t` switch, it is possible for the benchmark to complete, before the specified time period has elapsed, thus rendering the time variable irrelevant. Therefore, when using the `-t` and `-n` switches in tandem, it is important to consider whether your time variable exceeds the time it takes to process `n` requests.

```

macverse:~ estein$ ab -t 60 -c 10 http://localhost/~estein
This is ApacheBench, Version 1.3d <Revision: 1.73> $> apache-1.3
Copyright (c) 1996 Adam Twiss, Zeus Technology Ltd, http://www.zeustech.net/
Copyright (c) 1998-2002 The Apache Software Foundation, http://www.apache.org/

Benchmarking localhost (be patient)
Completed 5000 requests
Completed 10000 requests
Completed 15000 requests
Finished 16365 requests
Server Software:      Apache/1.3.33
Server Hostname:      localhost
Server Port:          80

Document Path:        /~estein
Document Length:      276 bytes

Concurrency Level:    10
Time taken for tests:  62.174 seconds
Complete requests:     16365
Failed requests:        0
Broken pipe errors:     0
Non-2xx responses:     16366
Total transferred:     7119218 bytes
HTML transferred:      4517816 bytes
Requests per second:   263.21 [#/sec] (mean)
Time per request:      37.99 [ms] (mean)
Time per request:      3.80 [ms] (mean, across all concurrent requests)
Transfer rate:         114.50 [Kbytes/sec] received

Connection Times (ms)
  min  mean[+/-sd] median max
Connect:  0    0  0.4   0   12
Processing: 0    2  1.0   2   19
Waiting:   0    2  1.1   2   19
Total:    0    3  1.2   2   19

Percentage of the requests served within a certain time (ms)
 50%  2
 66%  3
 75%  3
 80%  3
 90%  4
 95%  5
 98%  8
 99%  9
100% 19 (last request)
macverse:~ estein$

```

**Figure 2. Output of `ab -t 60 -c 10 http://localhost/~estein`**

## JMeter

JMeter is a pure Java application, developed under the purview of the Apache Jakarta project, which offers a wide array of tools for benchmarking all manner of client-serverbased applications. JMeter's modular architecture allows it to simulate a variety of load scenarios across multiple servers, networks, and objects. For the purposes of this article, however, we will only be using a subset of JMeter's capabilities to perform basic load testing. For information on using JMeter to perform detailed performance measures on more complex web applications you may refer to the user guide at <http://jakarta.apache.org/jmeter/usermanual/index.html>

Before we begin, it is instructive to go over the major architectural components that constitute JMeter.

- **Test Plan:** Defines the test sequence that JMeter will carry out during benchmarking.
- **Thread Group:** Used to delineate specific testing conditions, such as number of concurrent users, the Thread Group plays host to all other elements that comprise a test sequence.
- **Config Element:** Responsible for defining configuration information for the Sampler and Logic Controller elements.
- **Sampler:** A type of controller that designates both the protocol type and nature of requests within a given test plan.
- **Logic Controller:** A controller element that determines the sequence logic used by JMeter to initiate protocol requests.
- **Timer:** Used to manage the latency across queued user requests.
- **Listener:** Responsible for recording and visualizing of test sequence data.

## Getting Started

JMeter requires JDK 1.4 or higher to run properly. Type `java -version` in the Terminal application to discover which version of JDK is installed on your machine. Currently Apple supplies Java 1.3.1-1.4.2 Release 2 and J2SE 5.0 release 3 through regular software updates, as well as, at the following links:

<http://www.apple.com/downloads/macosx/apple/java131and142release2.html>

<http://www.apple.com/downloads/macosx/apple/java2se50release3.html>

In our tests, both JDK versions were able to run JMeter without a hitch. However, you will likely notice some performance improvement when using the J2SE 5.0 release 3. To specify this version as your system default or to change the precedence settings for your installed Java Virtual Machines (JVMs), use the Java Preferences utility located in `/Applications/Utilities/Java/Java2SE 5.0`.

The latest version of JMeter is available as a free download at [http://jakarta.apache.org/site/downloads/downloads\\_jmeter.cgi](http://jakarta.apache.org/site/downloads/downloads_jmeter.cgi). Depending on your preferences, you can download the application as source code or binary as either .zip or .tgz archives. The easiest option for most OS X users will be to download the binary in zip format, which can be easily uncompressed in the GUI by double clicking the `jakarta-jmeter-2.1.x.zip` archive. The resulting `jakarta-jmeter-2.1.1` folder can be placed in your Applications folder or other location within the filesystem. You can launch JMeter by either double clicking the `ApacheJMeter.jar` file, located in the `bin` directory of the `Jakarta-jmeter-2.1.1` folder, or by executing the `jmeter` shell script located in the same directory. Please note, that if you wish to run `jmeter` from the command line you will need to make the script executable using the `chmod` command. For command line interactive use, please refer to the “Getting Started”

section of the JMeter user manual: <http://jakarta.apache.org/jmeter/usermanual/get-started.html>




Figure 3. JMeter’s main window

## JMeter in Action

In this example, we will create a test plan that simulates 10 users initiating simultaneous requests to 2 different URLs and which repeats the test 100 times. Begin with the addition of a new Thread Group by right clicking (or control clicking) the **Test Plan** element in the left pane and selecting **Add > Thread Group**. Next, select the **Thread Group** element to display its configuration window and set the Number of Threads to 10 (representing users), the Ramp-Up Period to 0 (to initiate all users at once), and the Loop Count to 100.




Figure 4. Configuring Thread Group element

To add the next element of our test plan, HTTP Request Defaults, right click on the **Thread Group** element, and select **Add > Config Element > HTTP Request Defaults** (Figure 5).




Figure 5. Add HTTP Request Defaults

Select the HTTP Request Defaults icon under Thread Group and enter a server name (e.g. localhost or FQDN) in the configuration window (Figure 6).




Figure 6. Configuring HTTP Request Defaults element

Right click on the Thread Group element and then select Add>Sampler>HTTP Request twice, to create two HTTP Request elements, one for each web page. The Name and Path fields for each of the HTTP Request elements should be set to appropriate values. For the purpose of this example, we will name the first HTTP Request element home, and set its path to ~username/index.html and name the second HTTP Request another and set its path to ~username/another.html (figure 7).

**MACTECH**  
reseller network

**MacTech Helps  
You Profit!**

**Let the MacTech Reseller  
Network put you in the  
red-hot domain name  
registration and Internet  
business.**

**We've done the heavy  
lifting — product  
development, customer  
support, and  
infrastructure. All you  
need to do is  
sign up!**

[www.mactechdomains.com/resellers/](http://www.mactechdomains.com/resellers/)

**So take your  
business live on  
the Internet and start  
earning money TODAY!**




Figure 7. Configuring the “another” HTTP Request element

Now we need to add a timer to pause each request for a random period of time. Right click on the Thread Group and then select Add>Timer>Gaussian Random Time.

For the last element, we will add listeners by right clicking on the Thread Group element, then selecting Add>Listener>Aggregate Report. This particular listener lists, in tabular form, the response time, request count, min, max, average, error rate, and approximate number of request processes in seconds (e.g. Throughput), as well as, throughput in terms of Kilobytes per second. You may add additional listeners, such as Graph Results, View Results in Table, and View Results Tree, depending on your requirements.

To begin the test plan, simply select Run>Start from within the JMeter menu bar. You will be prompted to save the test plan, if you haven't already done so. During a successful run of your test plan, you may click on the **Aggregate Report** listener element (figure 8), or any other listener types you may have added, to view your dynamically updated test results. Figure 9, for example, demonstrates some of the graphing functions of JMeter's Graph Results listener element.




Figure 8. Aggregate Report listener view




Figure 9. Graph Results listener view

## Conclusion

In this article, we employed both ab and JMeter to “stress test” a locally hosted site via simple load scenarios. The amount of time a server takes to respond to clients’ requests across different load levels, is one of the most important metrics for evaluating web performance. Through the manipulation of load test variables, such as the number of concurrent requests, and with reference to tool-specific data visualization processes, we are better able to ascertain the extent to which these variables impact a server’s maximum sustainable load. Such information is critical for proper website management, planning and eventual performance tuning.

**MI**

## About The Authors

*Jin Lin and Emmanuel Stein are partners in the consulting firm MacVerse, Corp, which offers implementation, system administration, and development services geared towards the enterprise market. You may reach them at [info@macverse.com](mailto:info@macverse.com)*

**RADTECH™**  
Solutions that make sense.

**Gelz™**  
For iPod 4G, 5G, Nano & Mini

- Semi-transparent Silicone case
- Economical and colorful
- Good impact and drop protection

**Sleevz™**  
For iPod 3G, 4G, 5G, Nano & Mini

- Optex Super-80™ sleeve case for iPod
- The original control-through iPod Sleeve
- Includes transparent display protection

10% off your order with **Promo Code MT66Q**

[www.radtech.com](http://www.radtech.com)

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# Entourage 2004 Spotlight Support



An IT Perspective: How Microsoft Entourage 2004 now takes advantage of Spotlight

By Brian Johnson and Andy Ruff

## Introduction

In update 11.2.3, Microsoft added support for Spotlight and Sync Services to Microsoft Entourage 2004 running on Mac OS X 10.4. These two features allow users to search Entourage e-mail stores and to synchronize Entourage data with any software or hardware that takes advantage of sync services in the OS.

It's important for system administrators who need to plan deployment of this technology on Macs, and may need to consider configurations with many users, and with limited disk space, to understand how this all works. In this article, we'll focus on the Spotlight support added to Entourage. We'll tell you about how Spotlight support works in this update. Specifically, we'll also address Spotlight support considerations for multi-identity installations of Entourage.

## An Overview of Spotlight in Entourage

From the user perspective, Spotlight search in Entourage provides a mechanism that allows for the full text search of items in the Entourage database. Spotlight uses file based metadata and a constantly updating index to return results to queries passed through the Spotlight search interface in the operating system. Results return quickly because the index is updated based on messages coming from the file system. Once the initial indexing is complete on a set of data, additional data is indexed automatically as files update on the system.

One of the difficulties in making Spotlight work with Entourage had to do with how Entourage stores its data. All Entourage data is stored within a single database file per user identity. Entourage was designed to be multi-user at the application level. This was to allow multiple family members to have their own identities in the application in the home environment. Entourage stores its data in a single database for each user identity created in Entourage. When a user first sets up Entourage the Identity they get is named "Main Identity".

In order to support Spotlight searching, we had to develop a mechanism for providing Entourage's database content to Spotlight's file-oriented indexing process. We settled upon a solution that "mirrors" the essential item content and metadata to a series of cache files. As a new message arrives, we store the message within our database and spawn a cache file representing the message. When a user modifies a contact's phone number or changes the dates on an event, we update our database and the contents of each item's cache file. When Spotlight indexes Entourage, it is actually indexing the contents of each cache file rather than the Entourage database. This approach allows Spotlight's indexing process to work its magic on file change notifications, while not requiring a large overhaul of Entourage's data access architecture.

As an Entourage user's database potentially holds years of e-mail messages, the creation of cache files chances consuming large amounts of disk space for essentially redundant data. When we were considering this design, we found that through optimizations such as

writing only plain text content rather than HTML and ignoring e-mail attachments, we were able to generate a cache roughly 20% of the original Entourage database's size. We also decided that the feature would be optional, allowing any user to simply disable the creation of the cache within their Entourage preferences.




**Figure 1 – Spotlight is enabled by default for the first identity opened after update 11.2.3 is installed.**

The Spotlight preference pane in Entourage allows the user to both toggle the feature and rebuild the contents of the cache. On a moderately sized database of 200 MB, the creation of the cache file takes only a few minutes and happens in the background. The Rebuild button simply deletes all existing cache files, crawls the Entourage database, and generates a set of new cache files. A user would only need to rebuild if problems arise, as Entourage will continue to create, update, and delete cache files with each action performed on the Entourage database.

Once the cache files are created, Entourage's role in the indexing process is complete. Spotlight chooses when to index the cache files and how the results are displayed in the Spotlight Search Window, Smart Folders, and the Finder's Find functionality. As indexing progresses, the index in Spotlight is updated and queries containing the information the user is searching for begins to show up in the Search window. If you search for a set of words and Spotlight indexes an Entourage mail message with a matching phrase, the message will suddenly appear within the Spotlight Search Results Window. Figure 2 shows the results of a typical Spotlight search with Spotlight enabled in Entourage. The returned Entourage items can include mail, appointment, contact, task, and notes data.




**Figure 2: Spotlight search results with Entourage items returned.**

Double clicking on a returned item in Spotlight works as expected. You see the Entourage item open, just as if you had clicked on it in Entourage. So what's going on under the covers? Let's use some command line tools and take a look.

## Query with Command Line Tools

There are a number of command line utilities that we can use to query the Spotlight database. We can use these tools to see where Entourage is storing the Spotlight metadata that it's creating and we can also see what the metadata files themselves look like.

The first tool to look at is `mdfind`. `mdfind` queries the metadata store and returns the results of our query. This tool takes three parameters. The `-live` parameter will continuously scan the database for results and you'll see items added as they come into Entourage. The `-onlyin` parameter allows us to specify a particular folder for the search. Finally, the query parameter, a string representing the information that we're searching for. Apple's developer documentation provides more details on the syntax of Spotlight queries. Let's see if we can use this tool to find an Entourage item and see where the metadata is being stored:

Running the command "`mdfind welcomee@microsoft.com`" on my machine returns a result with the path:

/Users/Brianjo/Library/Caches/Metadata/Microsoft/Entourage/2004/Main Identity/Messages/OT/0B/0M/0K/1.vRgeMessage

The .vRgeMessage file is an Entourage mail message's cache file. When you perform a Spotlight search, the results always return cache files. As mentioned previously, cache files are merely file-based mirrors of Entourage database records with the metadata and content necessary for Spotlight indexing. The name of the cache file is the record ID for the corresponding database record. When a user opens the cache file from a Spotlight result, Entourage reads the filename, looks up the record ID within the database, and shows the item directly from the database.

The mdls command line utility allows you to see the metadata Spotlight has indexed for any given file. By passing the path to the 1.vRgeMessage cache file from our mdfind result to mdls, we can see Spotlight knows the following about the e-mail message:

/Users/Brianjo/Library/Caches/Metadata/Microsoft/Entourage/2004/Main Identity/Messages/OT/0B/0M/0K/1.vRgeMessage ———

```
com_microsoft_entourage_folderID      = 1
com_microsoft_entourage_messageSent   = 2006-03-21 00:23:21 -
0800
com_microsoft_entourage_recordID      = 1
com_microsoft_entourage_size          = 37783
kMDItemAttributeChangeDate            = 2006-03-21 21:55:25 -
0800
```

```
kMDItemAuthors                        = ("The Microsoft Mac
Team <WelcomeE@microsoft.com>")
kMDItemContentCreationDate            = 2006-03-21 00:23:21 -
0800
kMDItemContentModificationDate        = 2006-03-21 21:55:24 -
0800
kMDItemContentType                    =
"com.microsoft.entourage.virtual.message"
kMDItemContentTypeTree                = (
    "com.microsoft.entourage.virtual.message",
    "public.message",
    "public.data",
    "public.item"
)
kMDItemCoverage                       = "Inbox"
kMDItemDisplayName                    = "Welcome to Microsoft
Entourage 2004 for Macintosh"
kMDItemFSContentChangeDate            = 2006-03-21 21:55:24 -
0800
kMDItemFSCreationDate                 = 2006-03-21 21:55:24 -
0800
kMDItemFSCreatorCode                  = 0
kMDItemFSFinderFlags                  = 0
kMDItemFSInvisible                    = 0
kMDItemFSIsExtensionHidden            = 0
kMDItemFSLabel                        = 0
kMDItemFSName                         = "1.vRgeMessage"
kMDItemFSNodeCount                    = 0
kMDItemFSOwnerGroupID                 = 501
kMDItemFSOwnerUserID                  = 501
kMDItemFSSize                         = 6584
kMDItemFSTypeCode                     = 0
kMDItemID                             = 4306567
kMDItemKind                           = "Microsoft Entourage
message pointer"
kMDItemLastUsedDate                   = 2006-03-21 00:23:21 -
0800
kMDItemRecipients                     = ("New Microsoft
Entourage User ")
kMDItemTitle                          = "Welcome to Microsoft
Entourage 2004 for Macintosh"
```

# Save time... save money... synchronize your databases

## **fmSQL Synch eliminates double data-entry... forever.**

You can now synchronize your FileMaker and web (SQL) databases. Data is compared on a *record-by-record* as well as *field-by-field* basis. Existing databases can be configured for synchronizing with just a few modifications.

Some of the features include:

- ✓ Eliminates double data-entry
- ✓ Two-way synchronizing
- ✓ Field level conflict checking and resolution
- ✓ Setup Assistant
- ✓ Creates all SQL statements for you automatically
- ✓ Easily integrates with existing FileMaker solutions

Download a demo, and get more information at  
[www.fmSQL-Synch.com](http://www.fmSQL-Synch.com)



**Garrison Computer Services**

info@garrison.com.au

www.garrison.com.au

Ph: +61 2 4575-5247

PO Box 141 Windsor NSW 2756 Australia

# Stay in Control

wherever you go

with Timbuktu® Pro and Skype™, you can!



Choose how you want to manage your computer network and communicate with those working on or off-site, quickly, easily and securely—Mac or PC.

Windows

Mac OS X

## Timbuktu Pro version 8.6

- Skype Internet telephony integration
- Universal binary support for Macintosh

For IT professionals, telecommuters and those on the go, Timbuktu Pro is the award-winning standard in remote control technology for access to remote Windows and Macintosh desktops on LANs, WANs, the Internet, or through Skype telephony connections. For nearly 20 years, Timbuktu Pro has made it simple to find, connect, and control remote client machines from anywhere!

**Download a free evaluation at:**  
**[www.timbuktopro.com](http://www.timbuktopro.com)**



netopia®

BROADBAND WITHOUT BOUNDARIES®

Timbuktu Pro 8.6 works with Skype software. This product uses the Skype API but is not endorsed or certified by Skype. Skype is a trademark of Skype Technologies S.A. in Luxembourg and other countries.

© 2006, Netopia, Inc. All rights reserved. Netopia, the Netopia design, and Timbuktu are registered trademarks belonging to Netopia, Inc., registered in the U.S. Patent and Trademark Office. All other trademarks are the property of their respective owners.

Whenever possible, metadata provided by Entourage is designed so that attribute names and values match those used by an analogous Apple application (e.g. message title). We hope that this will allow anyone who builds a solution on top of Spotlight may easily support Entourage alongside Apple's applications. We only deviated by adding additional attributes – nearly all properties available in our AppleScript dictionary are available as metadata. The design is intended so that scripters can use Spotlight as a quick way to query information and round-trip interactions with results through AppleScript. It also makes it possible to create useful queries such as “all unread messages today”( `com_microsoft_entourage_unread == 1 && kMDItemContentCreationDate >= $time.today` ).

## Importing Metadata

Spotlight is designed such that it does not need to know about the file format of each file in order to index the file's contents. Instead, a developer provides a plug-in for Spotlight that handles both reading a file and returning metadata to the Spotlight indexing engine. These plugins are known as Metadata Importers and may be found `\Library\Spotlight` or are sometimes located within an application's bundle.




**Figure 3: The Microsoft Entourage.mdipl importer plug-in**

The metadata importer plug-in that's used with Entourage is called `Microsoft Entourage.mdipl importer`. When Spotlight comes across an Entourage-generated cache file, Spotlight passes the path of the file to the Entourage metadata importer, the importer reads the file, and then passes the metadata back to Spotlight. You can see the info for this plug-in in Figure 3. Notice that this plug-in is a universal binary and that it runs natively on an Intel-based Mac.

## Multiple Identities?

There are a few things that system administrators should understand if they are

going to use Spotlight search in multi-user scenarios. Given that Entourage can work as a multi-identity application, one thing you'll probably wonder about is, how does Spotlight know about the currently active identity in Entourage? The answer is that it doesn't. While we only automatically enable Spotlight indexing for the first identity launched after the update is applied, a user may turn on indexing of additional identities by enabling the Spotlight preference in Entourage. If the user then double clicks on an item, that item is only opened if its associated identity is currently active. Entourage actually uses the folder path to determine the identity of a result. If the identity is not currently active then the user will get the message shown in Figure 4.



**Figure 4. Trying to open an item associated with an inactive identity**

As you can imagine, in scenarios where many people use the same account on a Mac and then differentiate identities in Entourage, Spotlight could become pretty useless when trying to find specific e-mail items as search results intermix results across multiple identities. For that reason, Microsoft recommends that in situations where multiple users will want to use Spotlight search with Entourage, users should have their own user accounts set up on the Macs.

## Removing Spotlight from Entourage

There are a number of reasons a system administrator might want to completely disable Spotlight searching in Entourage. First, multiple user accounts on a machine are not always practical. In some cases, schoolrooms use a single account per classroom and kids are able to check their e-mail by simply switching identities in Entourage. On a machine with many dozens of identities, using Spotlight to find anything could be pretty difficult. Second, the cache used for Entourage content does take disk space. In a scenario where a user has a large entourage database, or there are multiple accounts on the machine with large databases, disk space can potentially become an issue. Finally there are privacy considerations around using Spotlight searches on Entourage content, especially if multiple identities are used on the same user account. Even if the searcher can't see the e-mail that's returned, they might be able to get more information than the user wants them to have about items returned in a search.


To completely disable Spotlight in Entourage, simply remove the `Microsoft Entourage.mdipl importer` plug-in

THINGS MADE IN CHINA

COMPRESSED BY

**STUFFIT DELUXE**

SOFTWARE FOR MAC/PC



from the /Library/Spotlight folder and restart Entourage. When you look at the Entourage preferences you'll see as in Figure 5, that the Spotlight preference is no longer available.




Figure 5: Spotlight removed from Entourage preferences

Finally, if you completely disable Spotlight in Entourage, remember to go to the /Users/<username>/Library/Caches/Metadata/Microsoft/Entourage/2004/ folder and delete any folders there that you no longer want to be available in Spotlight searches.

In this article, we've described how Spotlight works with Entourage 2004. The Spotlight search functionality added in update 11.2.3 fundamentally changes the way that you can work with e-mail, contacts, calendar items, and notes in Entourage, allowing you to instantly find data you need using the tools built right into Tiger. This update makes Entourage a true, first class citizen in the OS and it makes working with Entourage data on the Mac easier than ever.

## Bibliography and References

- Apple Computer. "Working with Spotlight".  
<http://developer.apple.com/macosx/spotlight.html>.  
M. Amir Haque. "Spotlight Support in Entourage 2004". *The Entourage Blog*,  
<http://blogs.msdn.com/entourage/archive/2006/03/17/553801.aspx>



## About The Author

**Brian Johnson is the Microsoft Entourage Product Manager. You can contact Brian by e-mail at [brianjo@microsoft.com](mailto:brianjo@microsoft.com), or you can read his blog at <http://bufferoverrun.net>.**

**Andy Ruff is an Entourage Program Manager, and is the author of The Entourage Blog (<http://blogs.msdn.com/Entourage>).**

Is your computer running slow?  
Maybe it needs a boost.

Make your life a little faster and easier,  
**get better RAM.**

**betterRAM.com**<sup>sm</sup>

Name brand quality. Value pricing. Lifetime guarantee. World class support.

Toll Free: (800) 895-3493 • Outside US/Canada: 805-494-9797 • Fax: 805-494-9798

[www.betterram.com](http://www.betterram.com)



# Salvation Awaits You.



reddot design award  
winner 2006

# Convert Now.

Change the way you live, work and play. Convert to Razer Pro|Solutions™ for faster speed, more accurate control, better response and superb precision. In essence, increase your productivity. It is the Way to perfection. Born from a marriage of form and function, the ProClick™ v1.6 is a godsend.

Surrender to Salvation. **Click & Convert Now.**

Works with both  
**PC and Mac**

Find out why you should convert @  
[www.razerpro.com/convert](http://www.razerpro.com/convert)

Available at:

**Newegg.com**  
[www.newegg.com](http://www.newegg.com)

**Buy.com**  
[www.buy.com](http://www.buy.com)

**PCConnection.com**  
[www.pcconnection.com](http://www.pcconnection.com)

**PCMail**  
[www.pcmall.com](http://www.pcmall.com)

**TechDepot**  
[www.techdepot.com](http://www.techdepot.com)

**Did you find this article helpful?**  
**Imagine a year of help.**



MacTech is already read every month  
by tens of thousands of readers.  
Readers that represent the very heart  
and soul of the Mac community. Join  
the crowd and sign up risk free today!

**Get a special deal on MacTech:**

**[store.mactech.com/sampler](http://store.mactech.com/sampler)**

**MACTECH®**

**Missing an issue?**

Coming soon the *MacTech CD Vol. 1-22*,  
containing every issue of MacTech ever printed!

# Advertiser/Product Index

|                                    |                 |
|------------------------------------|-----------------|
| 4D, Inc.                           | 75              |
| Absoft Corporation                 | 95              |
| Aladdin Knowledge Systems, Inc.    | 31              |
| Allume Systems, Inc.               | 8               |
| Allume Systems, Inc.               | 47              |
| Allume Systems, Inc.               | 73              |
| Allume Systems, Inc.               | 127             |
| BetterRAM.com                      | 128             |
| Bönig und Kallenbach oHG           | 100             |
| Brad Sniderman                     | 88              |
| Brian Loomis                       | 104             |
| CRYPTOCARD Corporation             | 58              |
| Daystar Technology                 | 61              |
| Equilibrium                        | 60              |
| FileWave (USA), Inc.               | 57              |
| Garrison Computer Services         | 124             |
| Idea Storage Networks LLC          | 33              |
| IGC, Inc. / MaxEMail.com           | 105             |
| Intego, Inc.                       | 63              |
| Intel Corporation                  | 19              |
| InterSystems Corporation           | 87              |
| JCHS Media Pte Ltd. / Mobile Juice | 42              |
| Kerio Technologies Inc.            | 26              |
| LTA Projects                       | 28              |
| MacForge.net                       | 76              |
| MacResource Computers & Service    | 96              |
| MacScripter                        | 103             |
| MacSpeech, Inc.                    | 78              |
| MacTech Magazine                   | 109             |
| MARWARE, Inc.                      | 109             |
| Maxell                             | 16              |
| Meta Communications                | 89              |
| Microsoft                          | 37              |
| MOST Training and Consulting       | 90              |
| Netopia, Inc.                      | 125             |
| NetTeam Consulting                 | 41              |
| Now Software                       | 10              |
| OlympicControls Corp.              | 77              |
| OmniPilot Software, Inc.           | 43              |
| Opera Software ASA                 | 24              |
| Other World Computing              | 27, 53, 114-115 |
| Ovolab                             | 12              |
| Peachpit Press                     | 69              |
| Plantronics                        | 51              |
| PremiumSoft CyberTech Ltd.         | 99              |
| Protective Solutions Inc.          | 109             |
| Quantum Corporation                | 65              |
| QuickerTek                         | 106             |
| RadTech, LLC                       | 120             |
| Razer USA Ltd.                     | 129             |
| Seapine Software, Inc.             | 81              |
| Spiderworks                        | 35              |
| SubRosaSoft.com, Ltd.              | 85              |
| Tellurium Communications, Inc.     | 97              |
| VisiStat, Inc.                     | 49              |
| Windows IT Pro                     | 39              |
| WorldSync, Inc.                    | 71              |

|                                                         |                 |
|---------------------------------------------------------|-----------------|
| Absoft Compilers • Absoft Corporation                   | 95              |
| Accelerators/Upgrades • Daystar Technology              | 61              |
| Aquazone • Allume Systems, Inc.                         | 73              |
| Bags • LTA Projects                                     | 28              |
| BookEndz • OlympicControls Corp.                        | 77              |
| Caché • InterSystems Corporation                        | 87              |
| Check It • Allume Systems, Inc.                         | 8               |
| CopyCatX/FileSalvage • SubRosaSoft.com, Ltd.            | 85              |
| CRYPTO-Server • CRYPTOCARD Corporation                  | 58              |
| DeBabelizer • Equilibrium                               | 60              |
| Digital Storage Manager • Meta Communications           | 89              |
| FileWave • FileWave (USA), Inc.                         | 57              |
| fmSQL Synch • Garrison Computer Services                | 124             |
| HASP • Aladdin Knowledge Systems, Inc.                  | 31              |
| Hosted Store • Brian Loomis                             | 104             |
| iListen • MacSpeech, Inc.                               | 78              |
| Intel Compiler • Intel Corporation                      | 19              |
| Internet Cleanup • Allume Systems, Inc.                 | 47              |
| Kerio Server Software • Kerio Technologies Inc.         | 26              |
| Laptop and iPod Cases • MARWARE, Inc.                   | 109             |
| Lasso • OmniPilot Software, Inc.                        | 43              |
| Law Offices • Brad Sniderman                            | 88              |
| Mac HelpMate • MOST Training and Consulting             | 90              |
| MacResource Computers • MacResource Computers & Service | 96              |
| MacScripter.net • MacScripter                           | 103             |
| MacTech Magazine • MacTech Magazine                     | 109             |
| MarketBlast • 4D, Inc.                                  | 75              |
| maxemail.com • IGC, Inc. / MaxEMail.com                 | 105             |
| Microsoft Office • Microsoft                            | 37              |
| Mobile Juice • JCHS Media Pte Ltd. / Mobile Juice       | 42              |
| Navicat • PremiumSoft CyberTech Ltd.                    | 99              |
| NetTeam Server • NetTeam Consulting                     | 41              |
| Now Up-to-Date • Now Software                           | 10              |
| Open Source Directory • MacForge.net                    | 76              |
| Opera • Opera Software ASA                              | 24              |
| Other World Computing • Other World Computing           | 27, 53, 114-115 |
| Peachpit Press • Peachpit Press                         | 69              |
| Phlink • Ovolab                                         | 12              |
| PhonePipe • Tellurium Communications, Inc.              | 97              |
| Plantronics • Plantronics                               | 51              |
| PowerBook Accessories • QuickerTek                      | 106             |
| RadTech • RadTech, LLC                                  | 120             |
| RAM • BetterRAM.com                                     | 128             |
| Razer • Razer USA Ltd.                                  | 4               |
| Screen Protection • Protective Solutions Inc.           | 109             |
| SDLT Drive • Quantum Corporation                        | 65              |
| Seefile • Idea Storage Networks LLC                     | 33              |
| SERVICE USB • Bönig und Kallenbach oHG                  | 100             |
| SpiderWorks ebooks • Spiderworks                        | 35              |
| Stuffit • Allume Systems, Inc.                          | 127             |
| SyncDek • WorldSync, Inc.                               | 71              |
| Tapes • Maxell                                          | 16              |
| Test Track Pro • Seapine Software, Inc.                 | 81              |
| Timbuktu • Netopia, Inc.                                | 125             |
| VirusBarrier • Intego, Inc.                             | 63              |
| VisiStat • VisiStat, Inc.                               | 49              |
| Windows IT Pro • Windows IT Pro                         | 39              |

The index on this page is provided as a service to our readers. The publisher does not assume any liability for errors or omissions.